

COGNITIVE ROBOTICS

EDITED BY
ANGELO CANGELOSI
AND
MINORU ASADA



Cognitive Robotics

Intelligent Robotics and Autonomous Agents

Edited by Ronald C. Arkin

A complete list of the books in the Intelligent Robotics and Autonomous Agents series appears at the back of this book.

Cognitive Robotics

Edited by Angelo Cangelosi and Minoru Asada

**The MIT Press
Cambridge, Massachusetts
London, England**

© 2022 Angelo Cangelosi and Minoru Asada

This work is subject to a Creative Commons CC-BY-ND-NC license.

Subject to such license, all rights are reserved.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Times New Roman by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Cangelosi, Angelo, 1967– editor. | Asada, Minoru, editor.

Title: Cognitive robotics / edited by Angelo Cangelosi and Minoru Asada.

Other titles: Cognitive robotics (M.I.T. Press)

Description: Cambridge, Massachusetts : The MIT Press, [2022] | Series: Intelligent robotics and autonomous agents series | Includes bibliographical references and index.

Identifiers: LCCN 2021031320 | ISBN 9780262046831 (hardcover)

Subjects: LCSH: Autonomous robots.

Classification: LCC TJ211.35 .C628 2022 | DDC 629.8/92—dc23

LC record available at <https://lccn.loc.gov/2021031320>

To Stella and Virginia (AC)

To the memory of my beloved son Ryu, and to Jin and Yuko (MA)

Contents

	Preface	ix
	Acknowledgments	xi
I	DEFINITION AND APPROACHES	
1	What Is Cognitive Robotics?	3
	Angelo Cangelosi and Minoru Asada	
2	Neurorobotics: Neuroscience and Robots	19
	Tiffany J. Hwu and Jeffrey L. Krichmar	
3	Developmental Robotics	41
	Minoru Asada and Angelo Cangelosi	
4	Evolutionary Robotics	59
	Stefano Nolfi	
5	Swarm Robotics	77
	Mary Katherine Heinrich, Mostafa Wahby, Marco Dorigo, and Heiko Hamann	
6	Soft Robotics: A Developmental Approach	99
	Luca Scimeca and Fumiya Iida	
II	METHODS AND CONCEPTS	
7	Robot Platforms and Simulators	123
	Diego Ferigo, Alberto Parmiggiani, Elena Rampone, Vadim Tikhonoff, Silvio Traversaro, Daniele Pucci, and Lorenzo Natale	
8	Biomimetic Skin	145
	Markellos Ntagios, Oliver Ozioko, and Ravinder Dahiya	
9	Machine Learning for Cognitive Robotics	165
	Tetsuya Ogata, Kuniyuki Takahashi, Tatsuro Yamada, Shingo Murata, and Kazuma Sasaki	

10	Cognitive Architectures	191
	David Vernon	
11	Embodiment in Cognitive Science and Robotics	213
	Tom Ziemke	
12	Ethics of Robotics	231
	Vincent C. Müller	
III	BEHAVIORAL AND COGNITIVE CAPABILITIES	
13	Intrinsic Motivations for Open-Ended Learning	251
	Gianluca Baldassarre	
14	Principles of Cognitive Vision	271
	Yiannis Aloimonos and Giulio Sandini	
15	Cognitive Robot Navigation	295
	Jiru Wang, Jianxin Peng, Rui Yan, and Huajin Tang	
16	Cognitive Robot Manipulation	315
	Yiming Jiang and Chenguang Yang	
17	Cognitive Control for Decision and Human-Robot Collaboration	337
	Erwin Jose Lopez Pulgarin, Ute Leonards, and Guido Herrmann	
18	Social Cognition	361
	Yukie Nagai	
19	Human-Robot Interaction	379
	Tony Belpaeme	
20	Language and Communication	395
	Angelo Cangelosi and Tetsuya Ogata	
21	Knowledge Representation and Reasoning	413
	Michael Beetz	
22	Abstract Concepts	433
	Alessandro Di Nuovo	
23	Robots and Machine Consciousness	453
	Antonio Chella	
	Contributors	475
	Index	477

Preface

Truth is verified only by creation or invention.
—Gianbattista Vico

Artificial intelligence (AI), machine learning, and robotics have become household terms following recent significant advances in AI for various applications in health care, in banking, and on the web and in the testing of robots in nuclear-decommissioning sites, as social companions for children and older people, and, very recently, as potential technologies to manage infection risks in the COVID-19 era. Notwithstanding this significant progress and momentum and the overpromising, in some cases, of what robots endowed with AI algorithms can actually do, the challenge of building machines with humanlike behavioral, cognitive, and social capabilities is a daring enterprise.

What cognitive robotics offers is a novel and insightful way to address the bold challenges of building AI-powered intelligent robots by taking inspiration from the way natural cognitive systems (i.e., humans, animals, biological systems) develop intelligence by exploiting the full power of the interactions between their bodies and their brains, the physical and social environments in which they live, and their phylogenetic, developmental, and learning dynamics. This is consistent with Vico’s philosophical approach that “truth is verified only by creation or invention.” That is, by creating or inventing something new, such as designing a computational cognitive architecture to control a cognitive agent, or developing a machine-learning model of intrinsic motivation and consciousness capabilities in robots, or running experiments to test a robot’s capabilities to sense, plan, and act in the world, we can verify the validity of a scientific theory, hypothesis, or model.

The term and field of *cognitive robotics* have their origins in the 1990s, and it is somewhat surprising that over the last thirty years of research in this field, no comprehensive publication has covered the breadth and depth of cognitively inspired intelligent robotic systems. This is exactly the aim of this book: to provide the first comprehensive, state-of-the-art coverage of cognitive robotics research and of its definition, approaches, methods, and applications. We will set the scene in part I (“Definition and Approaches”) by providing a systematic definition of the term *cognitive robotics* and an overview of its historical developments. This part will also include a detailed discussion of the five main, seminal approaches to cognitive robotics: developmental, neuro-, evolutionary, swarm, and soft

robotics. Part II (“Methods and Concepts”) further expands the primary methodologies and concepts employed in this field. These range from the analysis of the most commonly used cognitive robotics platforms and robot simulators to the case of biomimetic skin as an example of a hardware-based approach to cognitive robots. Two further methodological chapters examine the use of machine-learning methods and of cognitive architectures. Additionally, we look at theoretical considerations in cognitive robots, such as embodiment and the ethical implications of robotics and AI. The final part, III (“Behavioral and Cognitive Capabilities”), comprises a set of chapters covering the broad spectrum of robotics models, experiments, and applications with regard to various behavioral and cognitive capabilities. This ranges from intrinsic motivation and perception to social cognition and language and up to robot consciousness issues. Each of these chapters will also explicitly discuss the psychology and neuroscience findings and principles that have inspired the cognitive robots’ models and experiments.

The target readership of this volume includes master’s and PhD students who want to learn about the concepts and methods in the field as well as researchers interested in specific cognitive robotics models and experiments. The book is written for an interdisciplinary audience, balancing technical details and examples for the computational reader as well as theoretical issues and high-level descriptions of robot experiments for the empirical sciences reader.

We hope the reader will enjoy learning about the beneficial connection between psychology and neuroscience findings on cognitive development and learning in humans and animals and the design of intelligent robots.

Angelo Cangelosi and Minoru Asada

Acknowledgments

This volume is primarily the result of precious dedication on the part of the authors of the chapters. They are the pioneers in the field of cognitive robotics. Notwithstanding their busy jobs juggling teaching, research, paper and grant writing, and institutional administration roles, they kindly volunteered their time and effort in writing these chapters. Each author also acted as an anonymous cross-referee of other chapters, to assure the quality and clarity of the work.

We would also like to thank the staff at MIT Press—in particular Marie L. Lee for her enthusiasm and support for the book proposal and Elizabeth P. Swayze and Alex Hoopes for their support in the later stages of manuscript preparation. Special thanks go to Stella Cangelosi for the thorough help in the bibliography and manuscript formatting.

The initial idea of and work on this project began during one of Angelo Cangelosi's visits to the Artificial Intelligence Research Centre of the Japanese National Institute of Advanced Industrial Science and Technology (AIST). These visits were supported by an AIST grant to the University of Manchester. A particular thank you goes to the AIST director Junichi Tsujii, colleagues Tetsuya Ogata, Kristiina Jokinen, and Junpei (Joni) Zhong at AIST, and Sophia Ananiadou at Manchester for making this collaboration possible.

Angelo Cangelosi's work on cognitive robotics, as well as this volume, was also possible thanks to the generous support of research grants from the European Union Horizon 2020 program (i.e., the projects STRoNA, DCOMM, eLADDA, TRAINCREASE and PERSEO), the US Air Force Office of Scientific Research (project THRIVE++, AFOSR-EOARD Award FA9550-19-1-7002), the Honda Research Institute Europe (DeCIFER project), and the UK Research and Innovation Trustworthy Autonomous Systems Node in Trust.

Minoru Asada and his research group have been working in this area for more than a quarter century, and the following research projects contributed to their work, some of which are introduced in this book: Japan Science and Technology Agency Exploratory Research for Advanced Technology Asada Synergistic Intelligence Project (2005–2011); Japan Society for the Promotion of Science Grant-in-Aid for Specially Promoted Research on “Constructive Developmental Science Based on Understanding the Process from Neuro-Dynamics to Social Interaction” (2012–2016 PI: Minoru Asada); and Scientific Research on Innovative Areas, a Ministry of Education, Culture, Sports, Science and Technology Grant-in-Aid Project on “Constructive Developmental Science” (2012–2016 PI: Yasuo Kuniyoshi).

I DEFINITION AND APPROACHES

1 What Is Cognitive Robotics?

Angelo Cangelosi and Minoru Asada

1.1 Context and Definition

The wider field of robotics concerns the building of hardware mechatronics platforms with sensors and actuators to perform actions in the physical world and the designing of software solutions to link sensing and actuation in a purposeful—that is, intelligent—and adaptive way to achieve the task goal, with a variable degree of autonomy. This is captured, for example, in Matarić’s (2007, 2) definition of a robot as “an autonomous system which exists in the physical world, can sense its environment, and can act on it to achieve some goals.”

If we focus on the software side of robotics, the tools and approaches to building goal-oriented intelligent and adaptive capabilities in robots greatly overlap with the approaches and methods of artificial intelligence (AI). These range from good old-fashioned AI (GOF AI) knowledge-based reasoning and planning systems to the latest machine-learning algorithms of deep neural networks and reinforcement learning. Such a field combining robotics and AI can be referred to as “intelligent robotics” or, as recently proposed by Murphy (2019), “AI robotics.” Murphy (2019, 7) defines an intelligent robot as “a physically situated intelligent agent.” This designation is grounded in the concept of a robot being *physically situated* in the real world with an embodied physical structure suitable to perform a set task and the concept of an *intelligent agent* as a system that perceives its environment and takes actions to maximize its chances of success at adapting to the world. Such a definition and concepts of an intelligent robot practically coincide with Matarić’s general definition of a robot. In fact, the difference between (software) robotics and intelligent robotics is a really fuzzy distinction, as no researcher is really claiming to want to build “dumb” robots. Even the goal of modeling “Dumb Animals and Stupid Robots,” as Barbara Webb (1993) framed her project on the robot cricket, requires the use of nontrivial computer science and AI methods.

What is cognitive robotics then? Is it the same as intelligent robotics (AI robotics)?

Different definitions of cognitive robotics (CR hereafter) have been offered in the literature. In 1997 Stein proposed the first definition of CR when presenting the architectural principles for CR. Stein (1997, 471) defines CR as “the effort to build a physically embodied intelligent system—draws much of its approach from the cognitive sciences and natural

examples of embodied intelligent systems.” Kawamura and Browne (2009, 1) define CR as the “design and use of robots with humanlike intelligence in perception, motor control and high-level cognition,” stressing the need for interdisciplinary contributions from the various fields of robotics, AI, cognitive science, neuroscience, biology, philosophy, psychology, and cybernetics. Metta and Cangelosi (2012, 613) have proposed that CR is “the use of bio-inspired methods for the design of sensorimotor, cognitive, and social capabilities in autonomous robots.” All these definitions emphasize the role of an interdisciplinary approach to robot design and a focus on humanlike and bioinspired functions ranging from sensorimotor to higher-order cognitive functions, up to social skills. In particular, a fundamental influence in CR comes from the cognitive sciences, especially the disciplines interested in human cognition, such as psychology and neuroscience. This humanlike focus, however, does not exclude complementary insights from animal cognition and neuroscience in the design of bioinspired cognitive robots, such as tortoises and crickets (cf. Walter’s tortoises in sections 1.2 and 1.3.1).

Other researchers have characterized CR primarily as the distinctive focus on integrating higher-order functions, such as reasoning, to complement the standard intelligent robotics focus on sensing and action. De Giacomo (1998, 1), in the organization of the first meeting explicitly dedicated to CR (the 1998 Association for the Advancement of Artificial Intelligence [AAAI] Winter Symposium on Cognitive Robotics—see section 1.3.2), defined CR as the field “concerned with integrating reasoning, perception and action within a uniform theoretical and implementation framework.” Levesque and colleagues also focused on higher-order functions when defining CR as the “study of the knowledge representation and reasoning problems faced by an autonomous robot (or an agent) in a dynamic and incompletely known world” (Levesque and Lakemeyer 2008, 869; see also Levesque and Reiter 1998). As we will see in section 1.2, the emphasis on reasoning skills in the definition of CR is related to some of the influence of early AI knowledge representations experts in CR.

To summarize and integrate the various historical contributions to the characterization of CR, we would like to propose a comprehensive definition of CR that combines the above emphases on bioinspired—that is, humanlike and animallike—behavior and intelligence and on the distinctive interdisciplinary approach with strong contributions from the cognitive and neural sciences and from biology:

Cognitive robotics is the field that combines insights and methods from AI, as well as cognitive and biological sciences, to robotics.

Most of the current CR models typically focus on the design of one, or few, bioinspired sensorimotor and cognitive skills, as is the case in the CR models presented in part III of this volume. However, some works in CR also underscore the modeling of a system-level integration of a range of cognitive functions—for example, linking higher-level functions in reasoning and social skills with sensorimotor knowledge.

Now that we have defined CR, is this field the same as intelligent robotics (AI robotics)? In science it would be impossible, and counterproductive, to try to create an artificial, rigid distinction between different (sub)disciplines and approaches. Though one main distinction between CR and intelligent robotics lies in CR’s strong emphasis on designing bioinspired and cognitively inspired cognitive robots, in reality a continuum exists between the two fields. On one hand, there are CR models strictly constrained to known biological

mechanisms that are built to simulate and replicate the cognitive development phenomena observed in natural organisms. This is the case, for example, with Mori and Kuniyoshi's (2010) realistic rendering of the human fetus in their model of prenatal motor skill development (chapter 3) and of Morse et al.'s (2015) replication of child psychology experiments on the embodiment cues in early language learning (chapter 20). On the other hand, researchers have realized a variety of cognitive skills in intelligent robots via a combination of AI techniques without any justification for their biological inspiration or function.

The framing of CR as an integrative, systemic approach to modeling humanlike cognition in robots also explains its close link to the cognate modeling area of cognitive systems and its associated definition of cognition. The field of cognitive systems (a.k.a. artificial cognitive systems) refers to the creation of machines and software systems with humanlike cognition—that is, the “capacity for self-reliance, for being able to figure things out, for independent adaptive anticipatory action” (Vernon 2014, 2). Cognitive systems also tend to focus on higher-level cognition, on structured representations and systems perspectives, on influence from human cognition, and on exploratory research (Langley 2012). Cognitive systems as a discipline typically refers to the wider area of cognitive modeling with simulated and virtual agents, as well as physical robots, and to a variety of software-based agents and hardware-based smart objects (Morris et al. 2005; Vernon 2014). In its broadest sense, this has been extended to the design of intelligent human-computer interaction systems (a.k.a. cognitive systems engineering; Woods and Roth 1988) and to general-purpose AI systems such as the IBM Watson application (High 2012). With respect to CR, there is a good index of overlap when we consider the subareas of cognitive systems using physical robots, including cognitive systems of simulated robotics agents with a high degree of fidelity to the replication of body-environment physics dynamics.

Vernon (2014) considers four aspects when modeling artificial cognitive systems: 1) how much inspiration we take from natural systems, 2) how faithful we try to be in copying them, 3) how important we think the system's physical structure is, and 4) how we separate the identification of cognitive capability from the way we eventually decide to implement it. These aspects provide a method to position individual cognitive systems (and CR) models in a two-dimensional space where one axis defines the spectrum ranging from purely computational approaches to models strongly inspired by biological models, and the other axis defines the level of abstraction of the target biological model.

An important contribution from the field of cognitive systems is that of providing a more comprehensive operational definition of cognition. Following Vernon's (2014, 8) detailed characterization of cognition in artificial cognitive systems, cognition can be defined as “the process by which an autonomous system perceives its environment, learns from experience, anticipates the outcome of events, acts to pursue goals, and adapts to changing circumstances.” Thus, cognition can be seen as a systemwide process that integrates all of the capabilities of the agent within the key attributes of autonomy, perception, learning, anticipation, action, and adaptation. In particular, cognition can be represented as a cycle of anticipation, assimilation, and adaptation, embedded within a continuous process of action and perception and dynamically adapting via learning (figure 1.1).

This definition of cognition, and the identification of its six key attributes, can explain the variety of skills and capabilities the agent should possess: goal-oriented behavior,




Figure 1.1

The six key attributes of cognition in artificial cognitive systems. *Source:* Adapted from Vernon 2014.

autonomy, interaction via cooperation and communication, intention reading, interpretation of expected and unexpected events, prediction of the outcome of its own and of others' actions, action selection and evaluation, adaptation to changing circumstances, learning from experience, and monitoring and correcting its own performance (Vernon 2014).

This view of cognition is in line with the systemic and wider coverage of lower-level (perception and action) to higher-level (anticipation) capabilities of robots in CR. However, it places an emphasis on modeling the dynamic processes of cognition (assimilation, adaptation, learning). This is consistent with dynamical systems approaches in CR, such as in developmental robotics (cf. chapter 3).

The combined focus on the systemic and integrated approach to cognition, on the modeling of bioinspired humanlike and animallike cognitive capabilities, and on the interdisciplinary approaches to CR, as reflected in its definition above, will characterize the review of the state of the art in the chapters that follow. Of course, not all individual CR models aim to model the full breadth of behavioral and sociocognitive skills in a single robot. Typically, a specific CR model will implement a subset of such humanlike (and/or animallike) capabilities, depending on the specific task and skills the robot has to perform or the cognitive mechanisms the robot's model aims to operationalize and evaluate. This will be the case for most of the CR models and experiments presented in part III, with each chapter focusing primarily on a specific capability, from sensing, navigation, and manipulation to social and language skills to higher-level reasoning and consciousness.

Next we will look at the main epistemological and theoretical approaches to modeling behavior and intelligence that influenced and bootstrapped the emergence of the field of CR in the late 1990s. We will then summarize the origins and historical developments of CR.

1.2 Inspiration Principles and Theories

The early approaches to CR were influenced by both theoretical and computational stances in the modeling of behavior and cognition, in particular by the embodied cognition standpoint (e.g., Clark, Pfeifer) and by computational approaches to AI modeling of behavior-

based robotics and of higher-order reasoning function (e.g., Brooks). A further inspiration, particularly important from a historical point of view, was the direct influence of pioneering works on synthetic methods for modeling simple, animallike organisms (Walter, Braitenberg) and early computational neuroscience models for robotics (Edelman, Krichmar). Below we briefly discuss the specific theoretical and modeling works that motivated robotics researchers to take on the cognitive and bioinspired approach to intelligent robots and CR.

1.2.1 Embodied Cognition Theories

Embodied cognition is the approach to studying natural intelligent systems that underscores the roles of sensorimotor knowledge and representation and the interaction between our own body and the environment in producing intelligent behavior. In particular, the strong embodied cognition thesis states that the body plays a significant causal role, as a physically constitutive role, in the agent's cognitive processing (Wilson and Foglia 2017). A related approach is that of grounded cognition (Barsalou 2008; Pezzulo et al. 2013), which emphasizes the sensorimotor (“modal”) nature of the representations and internal simulation mechanisms (Vernon 2014). See chapter 11 for a detailed discussion on embodiment and embodied cognition.

Embodied cognition has affected various disciplines, including psychology (Pecher and Zwaan 2005; Barsalou 2008); cognitive sciences (Clark 1999); neuroscience (Pulvermüller and Fadiga 2010); and various computational modeling fields, such as language grounding (Cangelosi 2010), sensorimotor schema learning (Lara et al. 2018), and computational embodied neuroscience (Caligiore et al. 2010). Chapter 11 will also provide a detailed discussion of this issue and its specific contribution to CR.

In the very early stages of CR, there were two main theoretical stances on embodied cognition that have since been explicitly acknowledged to have influenced the very first cognitive robots. These are Andy Clark's (1999) theory on embodied cognitive science and Rolf Pfeifer's embodied intelligence and morphological computation stance.

Clark and Grush (1999) have specifically proposed a theoretical stance for a path toward CR. This is based on the “Cartesian agent” metaphor—that is, the combination of directly embodied, coupled, real-world action-taking with a decoupled, off-line reasoning capability. Thus, the cognitive phenomena of an agent involve off-line reasoning, which is vicarious environmental exploration and an internal representation.

This focus on the capability of having off-line reasoning functions grounded in embodied experience has had a strong impact on CR (Kawamura and Browne 2009) and has also contributed to some of the early CR emphasis on modeling knowledge representation and reasoning in robots (Levesque and Reiter 1998; Aiello et al. 2001).

This epistemological focus on higher-order cognition complements a parallel emphasis on the ability to develop cognition through sensorimotor coordination. This is the main stance proposed by Pfeifer and colleagues (Pfeifer and Scheier 2001; Pfeifer and Bongard 2006). Such an embodied cognition view is exemplified by the concept of “morphological computation”—that is, that certain sensorimotor and cognitive control processes are performed by the body and its interaction with the environment, rather than being performed by the brain. Pfeifer and Bongard (2006) use the example that the muscles and tendons of the human leg are elastic, and this directly influences locomotion control. When the leg impacts the ground while running, the knee performs small adaptive movements without

neural control. Thus, the control is supplied by the muscle-tendon system itself, which is part of the morphology of the agent. This morphological computation principle can also be exploited in robotics. A direct example of this is the “passive walker” (Collins et al. 2005; McGeer 1990), a simple robot that exploits gravity with a sloped track and the structure of two legs with flexible knees to move in a downward direction. This is possible without requiring any electric motors or electrical energy.

This attention to sensorimotor embodiment for cognition has greatly affected the development of CR, as many of the early cognitive robots have exploited the morphological computation principles (chapter 11). This is the case, for example, with soft robots exploiting the dynamics of the soft material of sensors and actuators (chapters 6 and 8), with evolutionary and swarm robotics for the automatic design of coupled body-brain-environment systems (chapters 4 and 5), and with developmental robotics and its application of the embodied cognition principles to motor development models (chapter 3).

1.2.2 AI and Knowledge-Based Systems

The classical (GOFAI) approach to AI, with its focus and breadth of methods for knowledge-based systems, symbolic representation, and reasoning, was also one of the key influences on CR. We have already mentioned early work by Levesque, Reiter, De Giacomo, and colleagues in the bootstrap of the CR discipline and community. In the 1998 AAAI Winter Symposium on Cognitive Robotics, many of the participants contributed to a “Cognitive Robotics Manifesto” with the explicit aim of modeling high-level robotic control in which robotic agents require reasoning using explicit knowledge representation systems that lead to a decision on how to act (Levesque and Reiter 1998; Aiello et al. 2001).

This approach follows the paradigm of perception-*reasoning*-action (or sense-*plan*-act), with a strong emphasis on the AI methods and models for reasoning/planning to connect robot sensing and action. It often involves the methods of situation calculus, description logic, and geometric reasoning typically applied to planning for action and navigation for the RoboCup challenge and mobile robot platforms (e.g., Woodbury and Oppenheim 1988; Aiello et al. 2001; but see Asada and von Stryk [2020] for a recent discussion of the scientific and technological challenges offered by the RoboCup challenge).

1.2.3 Behavior-Based Robotics

A different path to CR emerged from the alternative approach to AI based on the behavior-based robotics and the subsumption architecture proposed by Brooks (1991, 1996; Arkin 1998). In strong opposition to AI’s symbolic and representational methods, Brooks claims that intelligent behaviors can be achieved by reactive architectures, with a direct sense-act cycle and without the need for intermediate (symbolic) representations. This is exemplified by Brooks’s (1991) “Intelligence without Representation” nouvelle AI manifesto paper.

After the initial focus on mobile robot models of animal behavior (leading to the iRobot Roomba commercial vacuuming robot), the behavior-based robotics approach led to modeling behavior and cognition in humanoid robots (Brooks 1996; Matarić 1998). This included projects on the COG and the KISMET platforms (Brooks and Stein 1994; Brooks et al. 1998). This work explicitly led to an interdisciplinary approach using behavior-based

robotics as a tool for the synthesis of artificial behavior and the analysis of natural behavior, taking direct inspiration from cognitive science, neuroscience, and biology with methods from artificial life, evolutionary computation, and multiagent systems (Breazeal 2004). In the CR movement, this is closely linked to the development of evolutionary and swarm robotics (chapters 4 and 5, respectively).

1.2.4 Synthetic Methodologies

The “synthetic methodology” and “synthetic neural modeling” approaches to behavioral and cognitive modeling have also influenced CR (Krichmar 2012). These are methodologies based on the idea of recreating, in a simulated virtual environment or via physical platforms, embodied agents with a brain-inspired control system. They offer a balanced approach that emphasizes the intertwined interaction of the brain, the body, and the environment. The main synthetic methodologies directly influencing CR have come from Grey Walter’s “tortoises,” Valentino Braitenberg’s “vehicles,” and Chris Langton’s “artificial life” systems.

Grey Walter was a neuroscientist and pioneer in synthetic approaches to behavioral and cognitive modeling. In the late 1940s and early 1950s, he developed a set of electromechanical robots, called tortoises, capable of performing simple tasks such as phototaxis, following a light, and homing behavior, going to a battery-charging station. Walter’s first robot was called *Machina Speculatrix*, from the Latin verb *speculari*, which means “to explore,” as the tortoise actively explored the environment, as an animal would. Walter nicknamed two of the prototype robots ELSIE (from Electromechanical robot, Light Sensitive with Internal and External stability) and ELMER (ELectroMEchanical Robot; Walter 1950, 1953). He also proposed an electrical learning circuit named CORA (COnditioned Reflex Analogue) to model Pavlovian conditioning (Walter 1951). These systems implemented simple neural circuits. The focus on synthetic and neuroinspired modeling has galvanized many researchers in CR. For example, the Darwin series of robots developed by Edelman and colleagues (1992; Krichmar and Edelman 2003) follow on this synthetic methodology for mobile robots but with a stronger emphasis on using computational neuroscience models. This has led to the development of the CR neurorobotics approach (see chapter 2).

A subsequent synthetic modeling approach was proposed by the psychologist Braitenberg. In his well-known volume *Vehicles: Experiments in Synthetic Psychology*, Braitenberg (1986) describes a series of theoretical (fictional) models of simple mobile agents (i.e., vehicles). For example, Vehicle 1 is the simplest agent, with one sensor and one motor, and is capable of getting around by going straight with variable speeds depending on temperature sensors. Braitenberg describes a set of agents of increasing complexity in their sensorimotor system and the connectivity pattern between their sensors and motors and speculates on their ability to show behaviors that he describes as “fear and aggression” (Vehicle 2) and “love” (Vehicle 3).

These simple but elegant models of control in mobile agents have significantly influenced the field of CR, and of robotics and AI in general, as they provide an analysis of different control systems and their role in understanding behavior and cognition. For example, Hogg et al. (1991) developed a set of Braitenberg “creatures” as LEGO robots implementing and extending the various vehicles, and Hallam et al. (2002) used evolutionary computation to model the evolution of the spiking networks of Braitenberg’s controllers.

The third CR influential synthetic approach is that of artificial life (ALife; Langton 1997). This uses a prototypical synthetic methodology, as it aims to “synthetize” lifelike behavior and agents, in simulation and hardware. ALife models and applications go well beyond behavior and cognitive modeling; for example, they can be used to study artificial plants and artificial chemistry. In the early stage of ALife, significant emphasis was placed on agent and robot modeling, such as the CR evolutionary and swarm robotics approaches derived from building ALife agents (Steels and Brooks 1995). More recently, ALife has focused on synthetic biology and artificial chemistry, as well as on the origins of life.

1.3 History of Cognitive Robotics

Figure 1.2 gives a syncretic overview of the milestones in the history of CR, starting from the early attempts to model humanlike (and animallike) robots, which we call the “prehistory” of CR (from the early 1950s to the 1980s), to the period of the official start and establishment of the roots of CR (in the 1990s), to the contemporary evolution, diversification, and growth of various CR approaches (from 2000 onward). These historical developments will be discussed in detail.

1.3.1 Prehistory (1950–1980)

The tortoise robot models developed by Grey Walter (1950, 1953) in the early 1950s at the Burden Neurological Institute in Bristol, UK, can be considered the very first step in the (pre)history and origins of CR. Their novel synthetic methodology, the behavior-modeling focus, and the neuroinspired learning architecture pioneered by Walter have left a significant legacy not only in the field of CR but in the fields of robotics and AI in general (Holland 2003a, 2003b).

The 1960s saw the creation of the first intelligent robot, Shakey (Rosen et al. 1969; Nilsson 1984; see figure 1.2). It was developed between 1966 and 1972 at the Artificial Intelligence Center of the Stanford Research Institute (now SRI International). Shakey was a mobile robot capable of planning, route finding, and rearranging simple objects. The control architecture integrated sensing and action with the robot’s “model of the world.” This was implemented as a collection of predicate calculus statements in an indexed data structure, with five classes of entities (doors, wall faces, rooms, objects, robots) and a set of primitives to describe these entities in the model (e.g., distance between entities). For problem-solving, it used the QA3.5 theorem-proving system (Nilsson 1984). Shakey, and subsequent intelligent robots such as Flakey with its ability to follow and communicate with people, were the first platforms to experiment with linking AI with robotics, thus also influencing the AI robotics origins of CR.

The decade of the 1980s saw the creation of some of the seminal works that later influenced the development of CR. These include Braitenberg’s vehicles theoretical analysis and Brooks’s behavior-based robotics developments, as discussed in 1.2 (see figure 1.2). In the 1980s there is one work that, to the best of our knowledge, contains the first mention of the term “cognitive robotics.” This is the book *Principles and Elements of Thought Construction, Artificial Intelligence and Cognitive Robotics* by Charles Bowling (1987). It proposes a cognitive architecture for a simplified AI application based on the object calculus lattice (OCL) method.




Figure 1.2
The history of cognitive robotics with robot and book milestones.

1.3.2 Establishing Roots (1990s)

The first established gathering of a community explicitly using the title “Cognitive Robotics” and working at the interface of AI and robotics was the 1998 AAAI Fall Symposium on Cognitive Robotics (De Giacomo 1998). Giuseppe De Giacomo chaired it, with a strong presence from Ray Reiter’s team and their innovative work combining logic and reasoning capabilities in intelligent robots. In fact, this pioneering event helped roboticists to stress the higher-level cognitive functions of reasoning in action and perception robotic systems. It led to the “Cognitive Robotics Manifesto” (Levesque and Reiter 1998; Aiello et al. 2001). This event also provided the first definition of CR as the field “concerned with integrating reasoning, perception and action within a uniform theoretical and implementation framework” (De Giacomo 1998).

Other signs of the first attempts to focus on cognitively inspired robotics came from various groups working in AI and robotics, in addition to the work in behavior-based robotics and embodied cognition discussed above. In Japan, researchers working on cognitive skills design in humanoid robots started to define some of the principles of CR, such as exploring cognitive processes in systems with advanced cognitive functions by means of a “constructive approach” realized by repeating hypotheses and verification using robots (Asada et al. 1999).

1.3.3 Growth, Diversification, and Funding (2000s)

The CR roots established in the late 1990s, feeding from parallel contributions from the areas of behavior-based robotics, embodied cognition, and cognitive systems, led to a burst of growth in CR in the early to mid-2000s that still continues to this day. This is reflected by the flourishing workshops and special issues and seminal volumes in CR as well as further expansion of the associated CR approaches of developmental robotics, evolutionary robotics, and neurorobotics. For example, in 2002 leading pioneers in CR gathered in Bristol, UK, for the International Workshop on Biologically Inspired Robotics, dedicated to William Grey Walter (WGW02; Damper 2003; Holland 2003a). Another AAAI Winter Symposium on “The Intersection of Cognitive Science and Robotics: From Interfaces to Intelligence” was organized in 2004 (Shultz 2004). Other events included the 2006 Cognitive Robotics, Intelligence and Control Workshop (COGRIC) in Reading, UK (Becerra et al. 2006), the 2010 Dagstuhl Seminar “Cognitive Robotics” (Lakemeyer et al. 2010), and the 2013 international symposium in Osaka on “Past and Future Directions of Cognitive Developmental Robotics.”

This period also led to the diversification and growth of parallel, crosscutting CR approaches, each focusing on a specific learning or behavioral mechanism. These include developmental robotics, neurorobotics, evolutionary robotics, swarm/collective robotics, and soft robotics (as per part I of this volume).

The field of cognitive developmental robotics (Lungarella et al. 2003; Asada et al. 2009; Cangelosi and Schlesinger 2015) started in the early 2000s with the Workshop on Development and Learning (WDL; April 5–7, 2000, East Lansing, IL; cf. Weng et al. 2001) and the First International Workshop on “Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems” (EpiRob; September 17–19, 2001, Lund, Sweden; Zlatev and Balkenius 2001). The diffusion of baby robot platforms, such as the open systems iCub

robot (Metta et al. 2008, 2010) and the CB2 robot (Minato et al. 2007), significantly contributed to the growth of developmental robotics research (see figure 1.2 for the iCub and CB2 robots). See Cangelosi and Schlesinger (2015) and chapter 3 for a more recent and comprehensive review of the work in this field.

The field of neurorobotics is the subarea of CR that centers on the use of computational neuroscience and neuromorphic systems to control the robot's behavior and cognitive system (Browne et al. 2009; Krichmar 2012; see also chapter 2). This followed the early Darwin mobile robot models in the mid-1990s (Edelman et al. 1992) and led to numerous applications to mobile and humanoid robots, including the use of a neuromorphic system directly implementing hardware with neuron-like circuits (Rast et al. 2018) and the more recent neurorobotic platform in the Human Brain Project (Knoll and Gewaltig 2016).

Evolutionary robotics (Nolfi and Floreano 2000) is the CR approach to modeling the autonomous design of cognitive functions in robots via the use of evolutionary computation algorithms (see also chapter 4). This approach actually started in the mid-1990s, with subsequent growth in the 2000s along the wider evolutionary computation field and the CR/systems-oriented conference series “SAB: Simulation of Adaptive Behavior” and “ALIFE Artificial Life.” Evolutionary robotics benefited from the design and ease of access to small mobile robots in research laboratories, such as the Khepera robot (Mondada et al. 1999; see figure 1.2).

The field of swarm robotics can be seen as the application of swarm intelligence to robotics (see also chapter 5). This goes back to the early 1990s (e.g., Kube and Zhang 1992), with significant growth in the 2000s (e.g., Dorigo and Şahin 2004; Şahin 2004). The initial research in this field was mainly characterized by the transferring of biological principles, such as self-organization, to multirobot systems (Kube and Zhang 1992). Research in swarm robotics today generally focuses on specific methodologies, such as collective decision-making, as well as work toward applications—for example, for applications in sea monitoring, agriculture, and search and rescue.

More recently, the field of soft robotics has emerged as a branch of robotics, including CR, where soft and deformable materials are employed to endow robots with the ability to achieve more conformable, flexible, adaptable, and robust behaviors (Laschi et al. 2016; see chapter 6). This can also lead to the development of biomimetic (e.g., animal-inspired) robots such as octopus robots (Cianchetti et al. 2015; figure 1.2). This emphasizes concepts such as functional materials, deformable structures, and adaptive sensor morphology, which will be further discussed in chapter 6. The ability to devise and mimic unique, complex body dynamics and interactions with the physical world makes soft robots an exciting new field, where the limits of the (rigid) robots of the past century can be overcome for further understanding of bioinspired robotics and embodied cognition.

This period also saw interest and financial investment from various funding agencies worldwide in the growing areas of cognitive systems and CR. In 2002 the US Defense Advanced Research Projects Agency (DARPA) launched an initiative in cognitive systems to “develop the next generation of computational systems with radically new capabilities, ‘systems’ that know what they’re doing” (Brachman and Lemnios 2002).

The European Commission identified “Cognitive Systems” as one of the funding priorities for the new Sixth Framework Programme (FP7; 2002–2006), which then took on a more robotics-focused initiative with the “Cognitive Systems, Interaction and Robotics”

priority in the Seventh Framework Programme (FP7; 2007–2013; Maloney 2007). Examples of influential CR projects from these framework programs are RobotCub (which led to the iCub’s cognitive robot platform development; Metta et al. 2010; see also chapter 7; robotcub.org), CoSy for human-robot interaction using context-specific (situation and task) knowledge (Christensen et al. 2010), ITALK on developmental robotics for language grounding (Cangelosi et al. 2010), and POETICON/POETICON++ on the synthesis (poesis) of sensorimotor representations and natural language in everyday human interaction (Pastra 2008). This initiative also led to the funding of the network action grant EUCognition (www.eucognition.org/).

In 2003 the UK government’s Office of Science and Technology established a Foresight Project on “Cognitive Systems,” with subsequent interdisciplinary project funding from across the country’s different councils. This used the working definition of “Cognitive systems—natural and artificial—sense, act, think, feel, communicate, learn and evolve” (UK Foresight 2003; Morris et al. 2005). In this program, (cognitive) robotics was explicitly seen as a major example of one of the possible cognitive systems branches (along with computers, wearables, smart things, and so on).

In Japan, this led to the funding of large, collaborative projects in CR such as the Japan Science and Technology Agency Exploratory Research for Advanced Technology (JST ERATO) Asada Synergistic Intelligence Project and two Japan Society for the Promotion of Science (JSPS) Grants-in-Aid on “Constructive Developmental Science.”

1.4 Book Structure

This volume aims to provide a comprehensive, up-to-date overview of the state of the art in CR. As such, the chapters were authored by the leading international experts in the field, including many of the pioneers in CR.

In part I, we will first cover the main CR approaches or subareas—namely, neurorobotics, developmental robotics, evolutionary robotics, swarm robotics, and soft robotics.

Part II focuses on the methods and concepts common to most CR models and applications. It includes two chapters introducing the robot platforms and simulators and the bioinspired robot sensor and actuator technologies, a chapter providing an overview of machine-learning methods for CR, and two chapters on cognitive architectures and the concept of embodiment. It also contains a chapter on ethics for robotics, which is a fundamental concept in CR.

Part III is a series of chapters covering the whole spectrum of cognitive capabilities. Each chapter focuses on one specific behavioral/cognitive ability. Where appropriate, the chapter includes an explicit discussion of the bioinspired and cognitively inspired studies and theories that incited the subsequent robot models and experiments. This section of the book specifically includes chapters on the CR models of intrinsic motivation, visual perception, navigation and mapping, manipulation, human-robot interaction (HRI) decision and control, social cognition, human-robot interaction, language and communication, reasoning and knowledge representation, abstract concepts, and, finally, robot and machine consciousness.

This volume can be used to learn about the full breadth of approaches, methods, concepts, and models in CR—for example, for graduate students and researchers or as a reference book for a targeted effort on specific topics and work.

Each chapter also contains a section titled “Additional Reading and Resources” listing seminal papers and books in the specific topics covered by the authors, as well as links to internet and code resources. For general CR resources, see the “Introduction to Cognitive Robotics” course (www.cognitiverobotics.net). For pointers to software resources on CR, refer to the resources page of the Institute of Electrical and Electronics Engineers (IEEE) Technical Committee for Cognitive Robotics (<http://www.ieee-coro.org>).

References

- Aiello Luigia C., Daniele Nardi, and Fiara Pirri. 2001. “Case Studies in Cognitive Robotics.” In *Human and Machine Perception 3*, edited by V. Cantoni, V. Di Gesù, A. Setti, and D. Tegolo. Boston: Springer.
- Arkin, Ronald C. 1998. *Behavior-Based Robotics*. Cambridge, MA: MIT Press.
- Asada, Minoru, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. 2009. “Cognitive Developmental Robotics: A Survey.” *IEEE Transactions on Autonomous Mental Development* 1 (1): 12–34.
- Asada, Minoru, Hiroki Ishiguro, and Yasuo Kuniyoshi. 1999. “Toward Cognitive Robotics.” *Journal of the Robotics Society of Japan* 17 (1): 2–6.
- Asada, Minoru, and Oskar von Stryk. 2020. “Scientific and Technological Challenges in RoboCup.” *Annual Review of Control, Robotics, and Autonomous Systems* 3 (1): 441–471.
- Barsalou, Lawrence W. 2008. “Grounded Cognition.” *Annual Review of Psychology* 59:617–645.
- Becerra, Victor, Mark Bishop, William Browne, William Harwin, Kazuhiko Kawamura, Jeffrey L. Krichmar, Slawek Nasuto, Marica K. O’Malley, and Marjorie Skubic. 2006. *Cognitive Robotics, Intelligence and Control Workshop (COGRIC)*. Reading, UK. <http://www.cogric.reading.ac.uk>.
- Bowling, Charles M. 1987. *Principles and Elements of Thought Construction, Artificial Intelligence and Cognitive Robotics*. Csy Pub.
- Brachman R., and Z. Lemnios. 2002. *Darpa’s New Cognitive Systems Vision*. DARPA report. [http://www.defense-aerospace.com/articles-view/release/3/10422/darpa-research-into-cognitive-systems-\(june-17\).html](http://www.defense-aerospace.com/articles-view/release/3/10422/darpa-research-into-cognitive-systems-(june-17).html).
- Braitenberg, Valentino. 1986. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
- Breazeal, Cynthia L. 2004. *Designing Sociable Robots*. Cambridge, MA: MIT Press.
- Brooks, Rodney A. 1991. “Intelligence without Representation.” *Artificial Intelligence* 47 (1–3): 139–159.
- Brooks, Rodney A. 1996. “Behavior-Based Humanoid Robotics.” In Vol. 1, *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1–8. New York: IEEE.
- Brooks, Rodney A., Cynthia Breazeal, Matthew Marjanović, Brian Scassellati, and Matthew M. Williamson. 1998. “The Cog Project: Building a Humanoid Robot.” In *International Workshop on Computation for Metaphors, Analogy, and Agents*, 52–87. Berlin: Springer.
- Brooks, Rodney A., and Lynn Andrea Stein. 1994. “Building Brains for Bodies.” *Autonomous Robots* 1 (1): 7–25.
- Browne, William, Kazuhiko Kawamura, Jeffrey Krichmar, William Harwin, and Hiroaki Wagatsuma. 2009. Special issue, “Cognitive Robotics: New Insights into Robot and Human Intelligence by Reverse Engineering Brain Functions.” *IEEE Robotics and Automation Magazine* 16 (3): 17–18.
- Caligiore, Daniele, Anna M. Borghi, Domenico Parisi, and Gianluca Baldassarre. 2010. “TRoPICALS: A Computational Embodied Neuroscience Model of Compatibility Effects.” *Psychological Review* 117 (4): 1188.
- Cangelosi, Angelo. 2010. “Grounding Language in Action and Perception: From Cognitive Agents to Humanoid Robots.” *Physics of Life Reviews* 7 (2): 139–151.
- Cangelosi, Angelo, Giorgio Metta, Gerhard Sagerer, Stefano Nolfi, Chrystopher Nehaniv, Kerstin Fischer, Jun Tani, et al. 2010. “Integration of Action and Language Knowledge: A Roadmap for Developmental Robotics.” *IEEE Transactions on Autonomous Mental Development* 2 (3): 167–195.
- Cangelosi, Angelo, and Matthew Schlesinger. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Christensen, Henrik, Geert-Jan M. Kruijff, and Jeremy L. Wyatt, eds. *Cognitive Systems*. Vol. 8. Berlin: Springer.
- Cianchetti, M., M. Calisti, L. Margheri, M. Kuba, and C. Laschi. 2015. “Bioinspired Locomotion and Grasping in Water: The Soft Eight-Arm OCTOPUS Robot.” *Bioinspiration and Biomimetics* 10 (3): 035003.
- Clark, Andy. 1999. “An Embodied Cognitive Science?” *Trends in Cognitive Sciences* 3 (9): 345–351.

- Clark, Andy, and Rick Grush. 1999. "Towards a Cognitive Robotics." *Adaptive Behavior* 7 (1): 5–16.
- Collins, S., A. Ruina, R. Tedrake, and M. Wisse. 2005. "Efficient Bipedal Robots Based on Passive-Dynamic Walkers." *Science* 307 (5712): 1082–1085.
- Damper, Robert I. D. 2003. "WGW02. Proceedings of the International Workshop on Biologically Inspired Robotics, Dedicated to William Grey Walter." *Philosophical Transactions of the Royal Society A* 361 (1811): 2081–2421.
- De Giacomo, Giuseppe. 1998. "Cognitive Robotics." In *1998 AAAI Fall Symposium: Technical Report FS-98-02*. Menlo Park, CA: AAAI Press.
- Dorigo, Marco, and Erol Şahin. 2004. Special issue, "Swarm Robotics." *Autonomous Robots* 17 (2–3): 1–171.
- Edelman, Gerald M., George N. Reeke, W. Einar Gall, Giulio Tononi, Douglas Williams, and Olaf Sporns. 1992. "Synthetic Neural Modeling Applied to a Real-World Artifact." *Proceedings of the National Academy of Sciences* 89 (15): 7267–7271.
- Hallam, Bridget, Dario Floreano, Jean-Arcady Meyer, and Gillian Hayes. 2002. "Evolution of a Circuit of Spiking Neurons for Phototaxis in a Braitenberg Vehicle." In *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*, 335–344. Cambridge, MA: MIT Press.
- High, Rob. 2012. *The Era of Cognitive Systems: An Inside Look at IBM Watson and How It Works*. Armonk, NY: IBM Redbooks.
- Hogg, David Wardell, Fred Martin, and Mitchel Resnick. 1991. "Braitenberg Creatures." Cambridge, MA: Epistemology and Learning Group, MIT Media Laboratory.
- Holland, Owen. 2003a. "Exploration and High Adventure: The Legacy of Grey Walter." *Philosophical Transactions of the Royal Society of London A* 361 (1811): 2085–2121.
- Holland, Owen. 2003b. "The First Biologically Inspired Robots." *Robotica* 21 (4): 351–363.
- Kawamura, Kazuhiko, and Will Browne. 2009. "Cognitive Robotics." In *Encyclopedia of Complexity and Systems Science*, edited by R. Meyers. New York: Springer.
- Knoll, Alois, and Marc-Oliver Gewaltig. 2016. "Neurorobotics: A Strategic Pillar of the Human Brain Project." Supplement, Brain Inspired Intelligent Robotics, *Science Robotics* 354 (6318): 25–34.
- Krichmar, Jeffrey L. 2012. "Design Principles for Biologically Inspired Cognitive Robotics." *Biologically Inspired Cognitive Architectures* 1 (2012): 73–81.
- Krichmar, Jeffrey L., and Gerald M. Edelman. 2003. "Brain-Based Devices: Intelligent Systems Based on Principles of the Nervous System." In Vol. 1, *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Cat. No. 03CH37453, 940–945. New York: IEEE.
- Kube, C. Ronald, and Hong Zhang. 1992. "Collective Robotic Intelligence." In *Second International Conference on Simulation of Adaptive Behavior*, 460–468. Cambridge MA: MIT Press.
- Lakemeyer, G., H. J. Levesque, and F. Pirri, eds. 2010. "Cognitive Robotics." In *Dagstuhl Seminar Proceedings 10081*. <https://drops.dagstuhl.de/opus/portals/index.php?seminr=10081>.
- Langley, Pat. 2012. "The Cognitive Systems Paradigm." *Advances in Cognitive Systems* 1 (1): 3–13.
- Langton, Chris G., ed. 1997. *Artificial Life: An Overview*. Cambridge, MA: MIT Press.
- Lara, Bruno, Dadai Astorga, Emmanuel Mendoza-Bock, Manuel Pardo, Esaú Escobar, and Alejandra Ciria. 2018. "Embodied Cognitive Robotics and the Learning of Sensorimotor Schemes." *Adaptive Behavior* 26 (5): 225–238.
- Laschi, Cecilia, Barbara Mazzolai, and Matteo Cianchetti. 2016. "Soft Robotics: Technologies and Systems Pushing the Boundaries of Robot Abilities." *Science Robotics* 1 (1): eaah3690.
- Levesque, Hector, and Gerhard Lakemeyer. 2008. "Cognitive Robotics." *Foundations of Artificial Intelligence* 3:869–886.
- Levesque, Hector, and Ray Reiter. 1998. "High-Level Robotic Control: Beyond Planning, a Position Paper." In Vol. 37, *AIII 1998 Spring Symposium: Integrating Robotics Research: Taking the Next Big Leap*. N.p.
- Lungarella, Max, Giorgio Metta, Rolf Pfeifer, and Giulio Sandini. 2003. "Developmental Robotics: A Survey." *Connection Science* 15 (4): 151–190.
- Maloney, Colette. 2007. "The Commission Perspective: FP7 Challenge 2 Cognitive Systems, Interaction and Robotics." Paper presented at the EUcognition Project Meeting, January 12, 2007, Munich. http://www.vernon.eu/eucognition/six_monthly_meeting_2/Colette_Maloney.pdf.
- Matarić, Maja J. 1998. "Behavior-Based Robotics as a Tool for Synthesis of Artificial Behavior and Analysis of Natural Behavior." *Trends in Cognitive Sciences* 2 (3): 82–86.
- Matarić, Maja J. 2007. *The Robotics Primer*. Cambridge, MA: MIT Press.
- McGeer, Tad. 1990. "Passive Dynamic Walking." *International Journal of Robotic Research* 9 (2): 62–82.

- Metta, Giorgio, and Angelo Cangelosi. 2012. "Cognitive Robotics." In *Encyclopedia of the Sciences of Learning*, edited by N. M. Seel, 613–616. Boston: Springer.
- Metta, Giorgio, Lorenzo Natale, Francesco Nori, Giulio Sandini, David Vernon, Luciano Fadiga, Claes Von Hofsten, et al. 2010. "The iCub Humanoid Robot: An Open-Systems Platform for Research in Cognitive Development." *Neural Networks* 23 (8–9): 1125–1134.
- Metta, Giorgio, Giulio Sandini, David Vernon, Lorenzo Natale, and Francesco Nori. 2008. "The iCub Humanoid Robot: An Open Platform for Research in Embodied Cognition." In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, 50–56. Boston: Springer.
- Minato, Takashi, Yuichiro Yoshikawa, Tomoyuki Noda, Shuhei Ikemoto, Hiroshi Ishiguro, and Minoru Asada. 2007. "CB2: A Child Robot with Biomimetic Body for Cognitive Developmental Robotics." In *7th IEEE-RAS International Conference on Humanoid Robots*, 557–562. New York: IEEE.
- Mondada, Francesco, Edoardo Franzini, and Andre Guignard. 1999. "The Development of Khepera." In *Experiments with the Mini-Robot Khepera, Proceedings of the First International Khepera Workshop*, 7–14. Paderborn, Germany: Heinz Nixdorf Institute.
- Mori, Hiroki, and Yasuo Kuniyoshi. 2010. "A Human Fetus Development Simulation: Self-Organization of Behaviors through Tactile Sensation." In *2010 IEEE 9th International Conference on Development and Learning*, 82–87. New York: IEEE.
- Morris, Richard G. M., Lionel Tarassenko, and Michael Kenward. 2005. *Cognitive Systems-Information Processing Meets Brain Science*. San Diego: Elsevier.
- Morse, Anthony F., Viridian L. Benitez, Tony Belpaeme, Angelo Cangelosi, and Linda B. Smith. 2015. "Posture Affects How Robots and Infants Map Words to Objects." *PLoS One* 10 (3): e0116012.
- Murphy, Robin R. 2019. *Introduction to AI Robotics*. Cambridge, MA: MIT Press.
- Nilsson, Nils J. 1984. "Shakey the Robot." In *Technical Note 323*. Menlo Park, CA: AI Center, SRI International.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
- Pastra, Katerina. 2008. "PRAXICON: The Development of a Grounding Resource." In *Proceedings of the International Workshop on Human-Computer Conversation*. Bellagio, Italy.
- Pecher, Diane, and Rolf A. Zwaan, eds. 2005. *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*. Cambridge: Cambridge University Press.
- Pezzulo, Giovanni, Lawrence W. Barsalou, Angelo Cangelosi, Martin H. Fischer, Ken McRae, and Michael Spivey. 2013. "Computational Grounded Cognition: A New Alliance between Grounded Cognition and Computational Modeling." *Frontiers in Psychology* 3:612.
- Pfeifer, Rolf, and Josh Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, Rolf, and Christian Scheier. 2001. *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Pulvermüller, Friedemann, and Luciano Fadiga. 2010. "Active Perception: Sensorimotor Circuits as a Cortical Basis for Language." *Nature Reviews Neuroscience* 11 (5): 351–360.
- Rast, Alexander D., Samantha V. Adams, Simon Davidson, Sergio Davies, Michael Hopkins, Andrew Rowley, Alan Barry Stokes, Thomas Wennekers, Steve Furber, and Angelo Cangelosi. 2018. "Behavioral Learning in a Cognitive Neuromorphic Robot: An Integrative Approach." *IEEE Transactions on Neural Networks and Learning Systems* 29 (12): 6132–6144.
- Rosen, C. A., N. J. Nilsson, B. Rapahel, and R. O. Duda. 1969. *Research on Intelligent Automata*. SRI proposal. Menlo Park, CA: Stanford Research Institute.
- Şahin, Erol. 2004. "Swarm Robotics: From Sources of Inspiration to Domains of Application." In *International Workshop on Swarm Robotics*, 10–20. Berlin: Springer.
- Schultz, Alan. 2004. "The Intersection of Cognitive Science and Robotics: From Interfaces to Intelligence." In *2004 AAAI Fall Symposium: Technical Report FS-04-05*. Menlo Park, CA: AAAI Press.
- Steels, Luc, and Rodney Brooks, eds. 1995. *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*. London: Routledge.
- Stein, Lynn Andrea. 1997. "Postmodular Systems: Architectural Principles for Cognitive Robotics." *Cybernetics and Systems* 28 (6): 471–487.
- UK Foresight. 2003. *Foresight Report Setting out a Vision for the Future of Research in Natural and Artificial Cognitive Systems*. UK government report. <https://www.gov.uk/government/publications/cognitive-systems>.
- Vernon, David. 2014. *Artificial Cognitive Systems: A Primer*. Cambridge, MA: MIT Press.
- Walter, W. Grey. 1950. "An Imitation of Life." *Scientific American* 182 (5): 42–45.

- Walter, W. Grey. 1951. "A Machine That Learns." *Scientific American* 185 (2): 60–64.
- Walter, W. Grey. 1953. *The Living Brain*. London: Duckworth.
- Webb, Barbara. 1993. "Modeling Biological Behavior or 'Dumb Animals and Stupid Robots.'" In *Proceedings of the Second European Conference on Artificial Life*, 1090–1103. Edinburgh: University of Edinburgh, Department of Artificial Intelligence.
- Weng, Juyang, James McClelland, Alex Pentland, Olaf Sporns, Ida Stockman, Mriganka Sur, and Esther Thelen. 2001. "Autonomous Mental Development by Robots and Animals." *Science* 291 (5504): 599–600.
- Wilson, Robert A., and Lucia Foglia. 2017. "Embodied Cognition." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford, CA: Stanford University Press.
- Woodbury, Robert F., and Irving J. Oppenheim. 1988. "An Approach to Geometric Reasoning in Robotics." *IEEE Transactions on Aerospace and Electronic Systems* 24 (5): 630–646.
- Woods, David D., and Emilie M. Roth. 1988. "Cognitive Systems Engineering." In *Handbook of Human-Computer Interaction*, 3–43. Amsterdam: North-Holland/Elsevier.
- Zlatev, Jordan, and Christian Balkenius. 2001. "Why Epigenetic Robotics." Paper presented at the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems, September 17–19, Lund, Sweden.

2 **Neurorobotics: Neuroscience and Robots**

Tiffany J. Hwu and Jeffrey L. Krichmar

2.1 Introduction

Neurorobotics is the study of the interaction between neural systems and their physical embodiments on robotic platforms. Since the brain is strongly coupled with the body and situated within the surrounding environment, neurorobots can be a powerful tool for studying the intricate interactions between neural systems and the outside world. Neurorobotics also serves as a way to create autonomous systems that capture the advantages of biology for intelligent behavior. Compared to the general study of cognitive robotics, neurorobotics centers around biological brain functions—for example, the neural circuitry and functional anatomy that support basic cognitive processes. This chapter provides our viewpoints on this field, highlights some of its milestone events, and talks about its future potential.

2.2 Foundational Ideas in Neurorobotics

Many believe that neurorobotics got its beginning with Grey Walter's tortoises, which had simple light sensors and collision detectors attached to a basic analog circuit. His first robots, Elmer and Elsie, were programmed with simple reflexive neural circuits that controlled their movements based on the sensors. Despite the simplicity of these robots, complex and interesting behaviors emerged. For instance, one robot was placed in front of a mirror with a light on its nose. The robot started to react to its own presence in what could be interpreted as narcissistic behavior.

Braitenberg vehicles were another important example of complex behaviors emerging from simple circuitry. First introduced in the book titled *Vehicles* by Valentino Braitenberg (1986), a series of simple robots showed how basic neural circuits could create complex behaviors, some of which could even be attached to abstract human notions, such as emotion, with vehicle names like Fear, Aggression, Love, and Exploration. Each of these vehicles contained a light sensor and a motor on the left and right sides. In the vehicle displaying *fear*, the speed of each motor was directly proportional to the amount of light sensed by the sensor on the equivalent side. This caused the vehicle to speed away from the stimulus source, as if in fear. However, just crossing the wires caused the vehicle to speed toward the stimulus, as if in aggression. This simple robot provided an important

neuroscience lesson on the function of ipsilateral and contralateral connections in the nervous system. By making the motor speeds inversely proportional to the sensors, the vehicle displaying *fear* could turn into *love*, slowing down its movement toward the stimulus. Likewise, aggression then turned into exploration, gently seeking to be away from the stimulus. In this way, Braitenberg demonstrated how changing the balance of excitatory and inhibitory connections can affect behavior. Although the circuits themselves were simple, it was easy to place human interpretations on the resulting behaviors, teaching an important lesson that complex cognitive functions may actually be composed of very simple mechanics.

The Keck Machine Psychology Laboratory at the Neurosciences Institute in La Jolla, California, was also a source of foundational contributions in neurorobotics. Director Gerald Edelman (1987, 1993), whose work in immunology led to the Nobel Prize, advocated his theory of the nervous system in a book titled *Neural Darwinism: The Theory of Neuronal Group Selection*. The theory suggested there was selection of neural circuits during development through synaptic pruning and selection of groups of neurons during adulthood through reentrant connections. Important for neurorobotics was the notion of value systems to tie environmental signals to neuronal groups, which led to the selection of behaviors important for survival. As Edelman would say, “The brain is embodied, and the body is embedded in the environment.” Based on this idea, the group developed the Darwin series of Brain-Based Devices (Edelman et al. 1992; Reeke, Sporns, and Edelman 1990). Another phrase that drove this work was “The world is an unlabeled place,” which meant that perceptual categories must be selected through experience, rather than supervision. These Brain-Based Devices were robots with large-scale neural networks controlling their behavior (figure 2.1). However, these were not the feedforward-input neural networks that were popular then and became the deep neural networks of today. The Brain-Based Device’s neural networks contained anatomical details that resembled biological neural networks. There were sensory streams, top-down connections, and long-range connections between regions that were bidirectional as well as local lateral excitation and inhibition within brain regions. An early Brain-Based Device called Darwin V had an artificial nervous system that could learn preferences and predict the value of objects (Almassy, Edelman, and Sporns 1998). Although the robot was lumbering and did not exactly operate in real time, it did demonstrate operant conditioning and value-based learning.

One of the major venues in the early days of neurorobotics was the annual Simulation of Adaptive Behavior (SAB) conference. For example, SAB 2000 introduced a wide variety of exemplars, which would now be called neurorobots (Meyer et al. 2000). Arleo and Gerstner (2000) presented a model of head direction cells and hippocampal place cells, which was embodied on a Khepera robot, to demonstrate spatial navigation in the rodent. Arsenio (2000) created a neural circuit based on oscillators observed in the brain and showed how these could be used to realize humanoid arm movements and gait patterns. Collins and Wyeth (2000) introduced a cerebellar controller, based on Albus’s cerebellar model arithmetic computer (CMAC) neural network, to overcome delays when planning trajectories. Gonzalez and colleagues (2000) constructed a basal ganglia model to show action selection in a mobile robot. The robot would find cylinders, pick them up, and deposit the cylinders outside the wall of the robot arena. At this same meeting, Darwin VII, a Brain-Based Device capable of perceptual categorization, was introduced (Krichmar et al. 2000). For more details on Darwin VII, see the case study below. This is just a sampling of the work going on at this time.





Figure 2.1

Darwin IV Brain-Based Device. *Left:* Neural network model to control Darwin IV's behavior. *Right:* Darwin IV in a conditioning task. *Source:* Adapted with permission from Edelman et al. 1992.

The theme connecting the wide range of methods, robots, and behaviors at SAB 2000 was that neural network models were used to study some aspect of neuroscience by demonstrating behavior in a physical robot. Many of the researchers in these studies were pivotal in establishing the field of neurorobotics as it is known today.

Around this time period, other groups were creating robot designs that could be included within the field of neurorobotics. Rather than building brain circuits, they were investigating how the body and brain interact and how neural networks may develop. For example, Tony Prescott and his group at the University of Sheffield studied whisking in the rodent and developed a robotic sensorimotor circuit with biomimetic whiskers (Pearson et al. 2011). Figure 2.2 shows their Whiskerbot, which was completed around 2005. Dario Floreano helped establish the field of evolutionary robotics (Nolfi and Floreano 2000). Floreano and colleagues used evolutionary algorithms to evolve neural networks that supported a range of behaviors from navigating mazes to developing predator-prey strategies (Floreano and Keller 2010). For more details, the reader should refer to chapter 4. Rolf Pfeifer and Josh Bongard (2006) had the insight that the “body shapes the way we think.” They suggested that biological organisms perform morphological computation—that is, the body performs certain processes that would otherwise be performed by the brain.

Even though these biomimetic and evolutionary algorithms were not directly testing brain theories, they were increasing our knowledge of how the brain and body interact, and they were creating novel, biologically inspired algorithms and robot designs that would further the field of robots and AI.

As parallel-computing resources improved, some groups were approaching brain-scale neural simulations. Darwin VII’s neural network contained approximately twenty thousand neurons and nearly five hundred thousand synaptic connections, all of which had to be updated in real time to keep up with the active vision and sensors. The Darwin




Figure 2.2

Whiskerbot from the University of Sheffield. Whiskerbot had two active whiskers and a detailed neural network model to convert whisker deflection signals into simulated spike trains. *Source:* Adapted with permission from Pearson et al. 2011.

team used a Beowulf cluster with Message Passing Interface (MPI) to achieve real-time performance. Phil Goodman's Virtual Neurobot project had at least one hundred thousand highly detailed neurons on a computer cluster. Although the robot was virtual, it did need to respond in real time to recognize intent and trust in a human actor (Bray et al. 2012).

During this time there was often pushback from the community about the necessity for large-scale modeling. Many interesting results could be achieved with smaller neural networks, often with fewer than one hundred neurons. However, solving a problem in certain domains with small neural networks was unavoidable. For example, a model of the visual cortex that tested theories of feature binding and invariant object recognition (Seth et al. 2004b) required a neuron at every camera pixel (or receptive field) for each feature (two colors and four orientations). Since the network simulated the expansion of visual cortex receptive fields combining primitive features into objects (i.e., $V1 \rightarrow V2 \rightarrow V4 \rightarrow IT$), a large-scale neural network was necessary. However, applying the same modeling detail to a neural network that encoded tactile features with whiskers resulted in an order-of-magnitude-smaller network (Seth et al. 2004a).

In addition to practical reasons, large-scale modeling is often required to realize the neuronal dynamics and anatomical pathways observed in brain responses. Although this fidelity results in highly complex networks, it does allow one to test theories of the brain and make better predictions. Preserving anatomical projections leads to large-scale heterogeneous architectures. Having large groups of neurons with biophysical properties leads to interesting neural dynamics, as was observed in a large-scale model of the hippocampus and surrounding regions (Krichmar, Nitz, et al. 2005). In this model the complex interplay between the entorhinal cortex and the hippocampal subfields resulted in the reliance on different functional pathways at different points in the robot's learning (figure 2.3). Using large-scale neural models does come with a cost beyond computing power. At some point the neural network becomes so complex that it is as difficult to understand as the real brain. Interestingly, the analysis of the large-scale hippocampus model required the development of new tools; one was a recursive backtrace through neural activity (Krichmar, Nitz, et al. 2005), and the other applied Granger causality to the simulated neural network (Krichmar, Seth, et al. 2005).

Nowadays, large-scale neural network models are the norm. Neuromorphic hardware can support brain-scale neural networks at very low power (Indiveri et al. 2011; Merolla et al. 2014; Davies et al. 2018). Deep neural networks with many hidden layers are regularly developed (LeCun, Bengio, and Hinton 2015). With tools such as PyTorch and TensorFlow, graphics processing unit (GPU) clusters, and cloud computing, large-scale neural networks are within the reach of most researchers and students. Moreover, it turns out that size, in the form of many layers, is necessary to solve more challenging problems, such as image recognition (Krizhevsky, Sutskever, and Hinton 2017) or human-level game playing (Mnih et al. 2015).

2.2.1 Case Study: Darwin VII—Perceptual Categorization and Conditioning in a Brain-Based Device

Darwin VII was one of the first neurobots to demonstrate experience-dependent learning (i.e., learning by sampling the environment without supervisory signals) with a detailed,




Figure 2.3

Darwin X and a hippocampal model of episodic memory. (a) The overall neural network architecture included neuronal groups for the visual “what” and “where” streams ($V1 \rightarrow V2/4 \rightarrow IT$, $V1 \rightarrow V2/4 \rightarrow Pr$, respectively), head direction system (HD), reward system ($R+$, $R-$, S), and hippocampus. (b) Subfields within the hippocampus neural group. Arrows denote synaptic projections between subgroups. (c) Schematic of a dry variant of the Morris water maze. Colors denote landmarks; numbers denote starting positions of trials. (d) Darwin X Brain-Based Device. The hidden platform was a piece of black construction paper that Darwin X could not see with its camera but could detect with a downward-facing IR sensor. Adapted with permission from Krichmar, Nitz, et al. 2005.

neurobiologically plausible neural network (Krichmar and Edelman 2002). Darwin VII autonomously explored its environment and sampled stimuli that contained positive and negative values (figure 2.4). Through its experiences, Darwin VII built up perceptual categories of the objects it sampled. Darwin VII’s simulation was based on the anatomy and physiology of vertebrate nervous systems. The simulated nervous system comprised a number of areas labeled according to the analogous cortical and subcortical brain regions for vision, auditory processing, and value. Each area contained different types of neuronal units consisting of simulated local populations of neurons or neuronal groups. The simu-

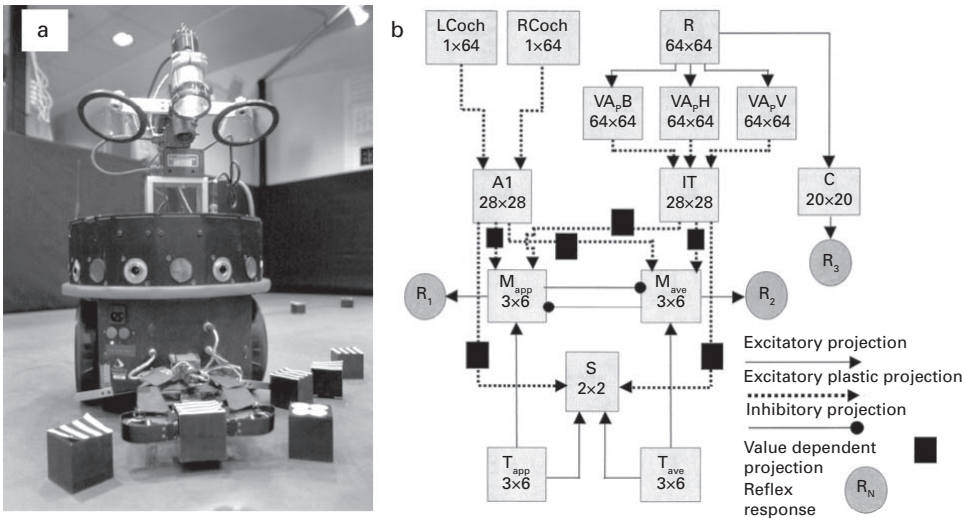


Figure 2.4

Darwin VII robot and neural network. (a) Darwin VII consists of a mobile base equipped with several sensors and effectors. Darwin VII is constructed on a circular platform with wheels that permit independent translational and rotational motion, with pan and tilt movement for its camera and microphones, and with object gripping by a one-degree-of-freedom manipulator or gripper. The CCD camera, two microphones on either side of the camera, and sensors embedded in the gripper that measure the surface conductivity of stimuli provide sensory input to the neural simulation. Eight infrared (IR) sensors are mounted at 45° intervals around the mobile platform. The IR sensors are responsive to the boundaries of the environment and were used to trigger reflexes for obstacle avoidance. All behavioral activity other than obstacle avoidance is triggered by signals received from the neural simulation. (b) The regional and functional neuroanatomy of Darwin VII. There are six major systems that make up the simulated nervous system: an auditory system, a visual system, a taste system, sets of motor neurons capable of triggering behavior, a visual tracking system, and a value system. The 64×64 gray-level pixel image captured by the CCD camera was relayed to a retinal area R and transmitted via topographic connections to a primary visual area VA_p . Three subdivisions in VA_p were selective for blob-like features, short horizontal line segments, or short vertical line segments. Responses within VA_p closely followed stimulus onset and projected nontopographically via activity-dependent plastic connections to a secondary visual area analogous to the inferotemporal cortex (IT). The frequency and amplitude information captured by Darwin VII's microphones was relayed to a simulated cochlear area (LCoCh and RCoCh) and transmitted via mapped tonotopic and activity-dependent plastic connections to a primary auditory area A1. A1 and IT contained local excitatory and inhibitory interactions producing firing patterns characterized by focal regions of excitation surrounded by inhibition. A1 and IT sent plastic projections to the value system S and to the motor areas M_{app} and M_{ave} . These two neuronal areas were capable of triggering two distinct behaviors, appetitive and aversive. The taste system (T_{app} and T_{ave}) consisted of two kinds of sensory units responsive to either the presence or absence of conductivity across the surface of stimulus objects as measured by sensors in Darwin VII's gripper. The taste system sent information to the motor areas (M_{app} and M_{ave}) and the value system (S). Area S projected diffusely with long-lasting, value-dependent activity to the auditory, visual, and motor behavior neurons. The visual tracking system controlled navigational movements, in particular the approach to objects identified by brightness contrast with respect to the background. To achieve tracking behavior, the retinal area R projected to area C ("colliculus"). *Source:* Adapted with permission from Krichmar and Edelman 2002.

lated nervous system contained 18 neuronal areas, 19,556 neuronal units, and approximately 450,000 synaptic connections. Figure 2.4b shows a high-level diagram of the different neural areas and the synaptic connections between neural areas in the simulated nervous system. A neuronal unit in Darwin VII was simulated with a mean firing-rate model, and the activity of such a unit corresponded roughly to the firing activity of a group of neurons averaged over a time period of 200 ms. This corresponded to the time needed to process sensory input, compute neuronal unit activities, update the connection strengths of plastic connections, and generate motor output.

The total contribution of synaptic input to unit i was given by

$$A_i(t) = \sum_{j=1}^N c_{ij} s_j(t)$$

where N is the number of connections to unit i , c_{ij} is the weight value of the connection projecting to unit i from unit j , and $s_j(t)$ is the activity of unit j at time step t . Negative values for c_{ij} corresponded to inhibitory connections. The activity level of unit i was given by

$$S_i(t+1) = \phi(\tanh(g_i(A_i(t) + \omega s_i(t))))$$

where

$$\phi_i(x) = \begin{cases} 0; & x < \sigma_i \\ x; & \text{otherwise} \end{cases}$$

and ω determined the persistence of unit activity from one cycle to the next, σ_i is a unit-specific firing threshold, and g_i is a scale factor, which differed depending on the neural area.

Connections within and between neuronal areas were subject to activity-dependent modification following a value-independent and a value-dependent synaptic rule. Synaptic modification was determined by both pre- and postsynaptic activity and resulted in either strengthening or weakening of the synaptic efficacy between two neuronal units. The Bienenstock, Cooper, and Munro (BCM) learning rule was used to govern synaptic change because it has a region in which weakly correlated inputs are depressed, and strongly correlated inputs are potentiated (Bienenstock, Cooper, and Munro 1982).

Value-independent synaptic changes in c_{ij} were given by

$$\Delta c_{ij}(t+1) = \varepsilon(c_{ij}(0) - c_{ij}(t)) + \eta s_j(t) F(s_i(t))$$

where $s_i(t)$ and $s_j(t)$ are activities of post- and presynaptic units, respectively, η is a fixed learning rate, ε is a decay constant, and $c_{ij}(0)$ is the initial ($t=0$) weight of connection c_{ij} . The decay constant ε governed a passive, uniform decay of synaptic weights to their original starting values. The function F is a piecewise linear approximation of the BCM learning rule.

The synaptic change for value-dependent synaptic plasticity was given by

$$\Delta c_{ij}(t+1) = \varepsilon(c_{ij}(0) - c_{ij}(t)) + \eta s_j(t) F(s_i(t)) \bar{S}$$

where \bar{S} is the average activity of the value system S (see figure 2.4b).

Darwin VII's environment consisted of an enclosed area with black walls and a floor covered with opaque black plastic panels, on which metallic cubes were distributed (figure 2.4a). The top surfaces of the blocks were covered with black-and-white patterns: blobs and stripes. Stripes on blocks in the gripper could be viewed in either a horizontal or vertical orientation, yielding a total of three stimulus classes of visual patterns to be discriminated (blob, horizontal, and vertical). A flashlight mounted on Darwin VII and aligned with its gripper caused the blocks, which contained a photodetector, to emit a beeping tone when Darwin VII was in the vicinity. The sides of the stimulus blocks were metallic and could be rendered either strongly conductive ("good taste," or appetitive) or weakly conductive ("bad taste," or aversive). Gripping of stimulus blocks activated the appropriate taste

neuronal units (either area T_{app} or area T_{ave}) to a level sufficient to drive the motor areas above a behavioral threshold. In the experiments, strongly conductive blocks with a striped pattern and a 3.9 kHz tone were arbitrarily chosen to be positive-value exemplars, whereas weakly conductive blocks with a blob pattern and a 3.3 kHz tone represented negative-value exemplars.

Early during the conditioning trials, Darwin VII picked up and “tasted” blocks that led to either appetitive or aversive responses (see figure 2.5a, *left panel*). During this period, it was the output of the taste neuronal units that activated the value system (S) and drove the motor neuronal units (M_{app} and M_{ave}) to cause a behavioral response. After conditioning, however, both the value system and the motor neuronal units were immediately activated upon the onset of IT’s response to a visual pattern or A1’s response to a tone. This shift from value system activity triggered in early trials by the unconditioned stimulus to value system activity triggered at the onset of the conditioned stimulus is analogous to the shift in dopaminergic neuronal activity found in the primate ventral tegmental area after conditioning (Schultz, Dayan, and Montague 1997).

After associating visual patterns with taste, Darwin VII continued to pick up and “taste” stripe-patterned blocks but avoided blob-patterned blocks (see figure 2.5a, *left panel*). After associating auditory sounds with taste, Darwin VII continued to pick up the high-frequency beeping blocks but avoided the low-frequency beeping blocks (see figure 2.5c, *left panel*). The right panel of figure 2.5b shows the percentage of conditioned responses, which were driven by the auditory or visual stimulus, for seven Darwin VII trials. The increase in conditioned responses showed that Darwin VII learned that auditory or visual cues predicted the value of the object, which resulted in it taking the appropriate behavioral response. These learning curves closely resembled those for similar conditioning experiments in rodents, pigeons, and other organisms.

In Darwin VII, activity in the simulated inferotemporal cortex, IT, provided the basis for visual perceptual categorization. Initially, IT’s responses to visual stimuli were weak and diffuse (see IT activity in figure 2.5a, *right panel*). After approximately five stimulus encounters, activity-dependent plasticity between primary visual cortex, VA_p , and IT caused IT responses to the different stimuli to become strong, sharp, and separable (see IT activity in figure 2.5b, *right panel*). Darwin VII’s object recognition was observed to be invariant with respect to scale, position, and rotation. Visual categorization of a stimulus occurred no matter where an object appeared in Darwin VII’s visual field, with the apparent size of the stimulus ranging from a maximum when the object was directly in front of Darwin VII to one-quarter of the maximum size when the object was distal to Darwin VII. Correct categorization of striped blocks in Darwin VII’s field of vision, when blocks were not in its gripper, occurred when the stripes on the blocks were rotated over a range of $\pm 30^\circ$ of a horizontal or vertical reference. These invariant category responses developed as a result of competition among activity-dependent plastic connections between retinotopically mapped VA_p and non-topographically mapped IT.

The behavior of Darwin VII showed that a robot operating on biological principles and without prespecified instructions could carry out perceptual categorization and conditioned responses. In both the perceptual categorization and conditioning experiments, the development of categorical responses required exploration of the environment and sensorimotor adaptation through specific and highly individual changes in connection strengths. Darwin VII




Figure 2.5

Left: Darwin VII during behavioral experiments. The panels to the right of Darwin VII show the activity of selected neural areas in the simulation (R, *top left*; IT, *top right*; A1, *bottom left*; M_{ave} , *bottom right, left side*; M_{app} , *bottom right, right side*). Each pixel in a selected neural area represents a neuronal unit, and activity is normalized in a range from no activity (*dark blue*) to maximal activity (*bright red*). (a) Darwin VII upon the first encounter with an aversive block. The stimulus block shown in this figure and in (b) had a blob-like visual pattern but did not beep. In this early conditioning trial, Darwin VII is shown picking up and “tasting” an aversive block. Activity in IT is insufficient, but activity in the taste system T_{ave} is sufficient to drive activity in the aversive motor behavior neural area (M_{ave}) above the behavioral threshold. (b) Darwin VII upon the tenth encounter with an aversive block having blob-like visual patterns. After primary conditioning with visual stimuli, activity in area IT is sufficient to drive the M_{ave} neuronal units above the behavioral threshold, triggering a motor response to avoid “tasting” an aversive block. (c) Darwin VII upon the tenth encounter with an aversive block having only auditory cues. After primary conditioning with auditory stimuli, activity in area A1 is sufficient to drive the M_{ave} neuronal units above the threshold to trigger a behavioral response. *Right:* The percentage of conditioned responses (%CR) per stimuli encountered by Darwin VII for auditory and visual stimuli. Each point is the average %CR for seven Darwin VII trials. *Source:* Adapted with permission from Krichmar and Edelman 2002.




Figure 2.5
(continued)

laid down groundwork for increasingly sophisticated neurorobots with more complex neural circuits and morphologies, which gave further insights into the relationships between brain, body, and behavior.

2.3 Building a Neurorobotics Community

Over the years, a neurorobotics community has emerged in part due to workshops and special journal issues on the topic. The *IEEE Robotics and Automation Magazine* devoted an issue to the topic (Browne et al. 2009). Special sessions were occasionally held on the topic at major IEEE robotics conferences. The European Union's Human Brain Project, a large-scale research project for understanding the nervous system, included a neurorobotics division headed up by Alois Knoll and Florian Rohrbach (Falotico et al. 2017).

In 2004, a special session on "Neurobotic Models in Neuroscience and Neuroinformatics" took place at the International Conference on the Simulation of Adaptive Behavior (Seth, Sporns, and Krichmar 2005). To introduce the session, it was stated that a neurobotic device has the following properties: 1) It engages in a behavioral task, 2) it is situated in a structured environment, and 3) its behavior is controlled by a simulated nervous system designed to reflect, at some level, the brain's architecture and dynamics. The session included Auke Ijspeert's research on evolving neural networks for a robotic salamander (Ijspeert, Crespi, and

Cabelguen 2005; Ijspeert et al. 2007). In this research, different motor patterns (i.e., swimming or walking) emerged due to the interaction between brain and body with the specific environment (i.e., water or land). Olaf Sporns and Max Lungarella showed how embodiment can alter and improve information processing in a neural system (Lungarella et al. 2005). In addition, several papers on how the hippocampus contributes to spatial memory were presented (Arleo, Smeraldi, and Gerstner 2004; Banquet et al. 2005; Chavarriaga et al. 2005; Krichmar, Seth, et al. 2005).

Robot models of rodent navigation have made up a number of neurorobotic implementations. One reason for the interest in these models is because robot navigation is a fascinating and complex problem. Another reason is that the neural activity patterns observed in the rat are clear, interesting, and amenable to modeling. For example, a head-direction cell can be modeled with an attractor network and cosine tuning curves (Stringer et al. 2002). A hippocampal place cell can be modeled with a two-dimensional Gaussian (Foster, Morris, and Dayan 2000). The more recent finding of grid cells in the entorhinal cortex has led to a number of proposed neural models (Zilli 2012). Using attractor networks and neural elements that resemble head direction cells, place cells, and grid cells, the Australian RatSLAM team has reported results with neuro-inspired algorithms that are as good as or better than state-of-the-art localization and mapping by conventional robots (Milford et al. 2016). Although great progress has been made in the conventional robotics community with SLAM, or simultaneous localization and mapping (Kohlbrecher et al. 2011; Mur-Artal, Montiel, and Tardos 2015) and path planning (LaValle 2011a, 2011b), a number of open issues still remain when it comes to flexible navigation under dynamic conditions. Under these challenging situations, rodents show superior performance and robustness and still provide inspiration for improved robot navigation algorithms.

2.4 Neurorobotics and Neuromorphic Engineering

An important potential development for the field of neurorobotics is the reemergence of neuromorphic engineering (Indiveri et al. 2011). By reemergence, we mean that the original analog circuits developed by Carver Mead (1990) and his team in the 1980s have led to near-commercially viable computers designed by large companies such as IBM (Merolla et al. 2014) and Intel (Davies et al. 2018). Like neurorobotics, neuromorphic engineering uses inspiration from the brain to build computer architectures and sensors. Because these computers were specifically designed for asynchronous, event-driven processing, spiking neural networks that controlled neurorobots were ideal for these platforms. Moreover, neuromorphic architectures hold great promise for neurorobot applications due to their low power budget and their fast, event-driven responses. For example, the SpiNNaker neuromorphic computer from Manchester has been used in an obstacle avoidance and random exploration task (Stewart et al. 2016). In addition to running neural networks on specialized hardware, very low power neuromorphic vision and auditory sensors are being developed (Liu and Delbruck 2010). Similar to biology, these sensors only respond to change or salient events, and when they do respond, it is with a train of spikes. This allows seamless integration of these sensors with spiking neural networks, and their event-driven nature leads to power efficiency that's ideal for embedded systems, such as robots.

The development of lightweight neuromorphic chips inspired the idea that many computing processes related to outdoor navigation could be implemented on neuromorphic hardware to control ground robots. Neuromorphic hardware is especially beneficial for outdoor navigation, as the robots must rely on battery power for long periods of time and are often used in vital operations such as search and rescue. Spiking implementations of low-level perceptual navigation tasks as well as high-level planning tasks allow for navigation subtasks to run in parallel.

Working with IBM's low-power TrueNorth neuromorphic chip (Esser et al. 2016), we demonstrated that a convolutional neural network (CNN) could be trained to self-drive a robot on a mountain trail (Hwu et al. 2017). Initially, the robot was driven along the trail using remote control. The RGB camera frames, along with the corresponding action controls of steering left, steering right, and driving forward, were recorded for training the CNN. The CNN was first trained with conventional backpropagation techniques, using the RGB images as input and the set of three actions as output. The weights of this neural network were then transferred to weights in a spiking neural network of the same structure as the original CNN. This spiking network was run on the TrueNorth chip, which was powered by the same single hobby-level nickel metal hydride (NiMH) battery used to power the motors of the robot (figure 2.6). The advantage of using this pipeline was that we were able to harness well-developed techniques of CNN training while achieving order-of-magnitude gains in energy efficiency. The circuit diagram and pipeline shown in figure 2.6 could generalize to other hardware and neurorobot applications.

2.4.1 Case Study: Spiking Wavefront Propagation—Brain-Inspired Neuromorphic Path Planning

Navigation is a necessary component of most robots and animals, both of which operate under the constraints of limited time and energy. Using inspiration from brain connectivity, neuron spiking dynamics, and a recent finding that axonal conductance undergoes experience-dependent plasticity (Fields 2015), a model of spiking wavefront propagation was created (Hwu et al. 2018). The model was inspired by the role of the hippocampus in animal navigation. This includes the existence of place cells in the hippocampus, which are active according to the physical location of the animal (O'Keefe and Dostrovsky 1971). These place cells are involved in hippocampal replay, in which the place cells activate in sequence according to potential trajectory routes the animal can take (Dragoi and Tonegawa 2011; Pfeiffer and Foster 2013). Another biological observation behind spiking wavefront propagation is that spreading waves of activity can be found across several areas of the brain including the hippocampus, supporting brain connectivity and memory (Zhang and Jacobs 2015).

Combining these observations, the model of spiking wavefront propagation is able to plan paths through a grid representation of space. Each grid unit corresponds to a discretized area of physical space, and connections between units represent the ability to travel from one area to a neighboring area. Each unit in the grid represents a single neuron with spiking dynamics. The membrane potential of neuron i at time $t + 1$ is represented by

$$v_i(t + 1) = u_i(t) + I_i(t),$$

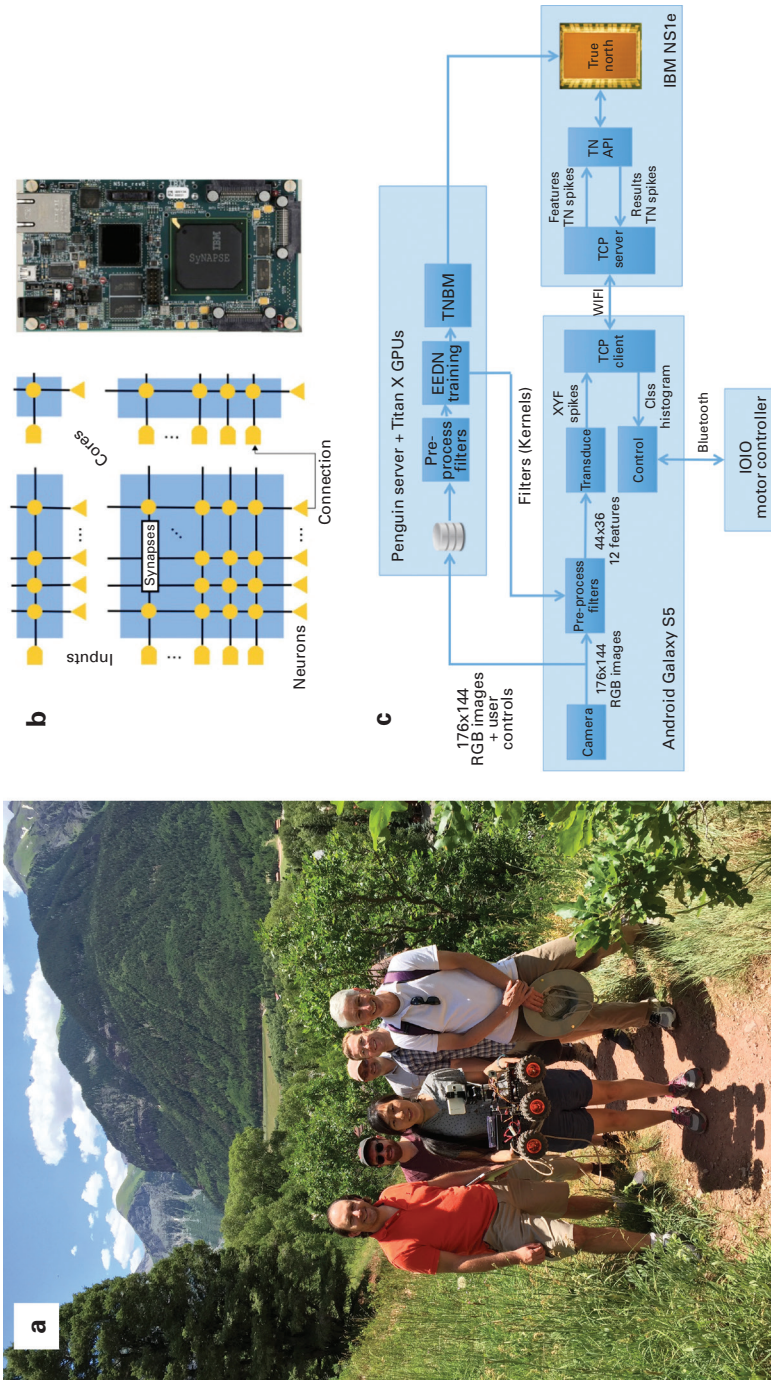


Figure 2.6

A self-driving robot using deep convolutional neural networks on IBM's TrueNorth neuromorphic hardware. *Left to right:* Rodrigo Alvarez-Icaza (IBM), Jacob Isbell (University of Maryland), Tiffany Hwu (University of California, Irvine), Will Browne (Victoria University of Wellington), Andrew Cassidy (IBM), and Jeff Krichmar (University of California, Irvine). Missing from the photograph is Nicolas Oros (BrainChip). *(b, Left):* The connectivity of the IBM TrueNorth neuromorphic chip. *Right:* An image of the IBM TrueNorth NS1e board used in the experiments. *(c)* Data pipeline for running the self-driving robot. Training was done separately with the Eedn MatConvNet package using Titan X GPUs. During testing, a Wi-Fi connection between the Android Galaxy S5 and IBM NS1e transmitted spiking data back and forth, using the TrueNorth (TN) runtime API. *Source:* Adapted with permission from Hwu et al. 2018.

in which $u_i(t)$ is the recovery variable, and $I_i(t)$ is the input current at time t . The recovery variable $u_i(t+1)$ is modeled as

$$u_i(t+1) = (-5 \text{ if } v_i(t) = 1; \min(u_i(t) + 1, 0) \text{ otherwise}),$$

such that if it starts as a negative value, it increases at a steady rate toward a baseline value of 0. The input current $I_i(t+1)$ is represented as

$$I_i(t+1) = \sum_j (1 \text{ if } d_{ij}(t) = 1; 0 \text{ otherwise}),$$

such that $d_{ij}(t)$ is the delay counter of the signal from neighboring neuron j to neuron i . The delay $d_{ij}(t+1)$ is calculated as

$$d_{ij}(t+1) = (D_{ij}(t) \text{ if } v_j(t) = 1; \max(d_{ij}(t) - 1, 0) \text{ otherwise}),$$

such that it behaves as a timer corresponding to axonal delay with a starting value of $D_{ij}(t)$. This starting value of $D_{ij}(t)$ is a delay value depending on the cost of traversing the spatial area corresponding to the neuron. Taken together, these equations describe the simplified dynamics of a spiking neuron. When a spike from a neighboring neuron occurs, the input current I_i is set to 1, causing a spike. Immediately after, the recovery variable u_i is set to -5 , which then counts up by 1 at each successive time step and stops at 0. This mechanism models the refractory period of the neuron. Next, all delay counters d_{ij} for all neighbor neurons j are set to their assigned starting values of D_{ij} .

Multiple possibilities exist for encoding the values D_{ij} . These values should encode the cost of traversing from one area to another. This may be the energy required, the potential risks, or the physical wear. For instance, traveling through rough terrain would be riskier and require more energy for ground robots and therefore have higher costs. A cost map of the same dimensions as the grid can transfer to values of D_{ij} in a one-to-one fashion. The cost map, if known in advance, can be used to populate delay values of the grid prior to running spiking wavefront propagation. They may also be learned on the fly while exploring the terrain. In neuroscience, this would correlate to axonal plasticity, in which the myelin sheath of a neuron consisting of white matter grows in volume with heightened activity and subsequently increases the speed of signals traveling from one neuron to another (Fields 2015). As an agent travels through its environment, either randomly or by intentionally navigating, D_{ij} values are updated each time the agent enters a new grid area using the following equation:

$$D_{ij}(t+1) = D_{ij}(t) + \delta(\text{map}_{xy} - D_{ij}(t)),$$

where δ is the learning rate, and map_{xy} is a sample of the cost as the agent traversed location coordinates (x, y) corresponding with grid neuron i . The update rule is applied for each of the j neighbors of neuron i . The advantages of axonal plasticity are that the agent can learn while operating, continuously gaining new information. With a small learning rate, the model accounts for noise in the environment such that if the agent samples a faulty cost value due to sensor error or environmental factors, the effect is averaged across multiple trials. However, learning accurate cost values for an entire grid may require many trials, as each grid area must be traversed several times. It may therefore be preferable to start with an a priori map of costs, updating with sensor-based observations as they occur.

To perform path planning using the grid encoded with costs, an input current is added to the neuron corresponding to the location of the agent to induce a spike. This induces spikes in neighboring neurons, subsequently starting a traveling wave across the entire grid. As the spikes occur, their spikes are recorded using address event representation (AER), which includes pairs of neuron IDs and spike times. Figure 2.7 shows how using AER can be used for path planning.

To plan a path from the start location to any other location, the first spike time of the destination neuron is recorded. The ID of the destination neuron is also recorded on a list. Then the spike times of each neighboring neuron are examined, and the neuron with the most recent spike is appended to the list. The same process is repeated with this neighboring neuron, and so on, until the start neuron is added to the list. The optimal path accounting for length and cost is then returned as the reversed list of neuron IDs.

The present spiking wavefront algorithm was successfully tested on a mobile ground robot traversing over grass, dirt, and asphalt terrains (Hwu et al. 2018). The robot was created from affordable hobbyist parts and an Android phone for computation (figure 2.8, *bottom left*). The robot motors and sensors were powered by a single NiMH battery, making energy savings a priority in its operation. The robot was tested at a large outdoor park in two areas (figure 2.8, *top left*). One area was a grass field surrounded by an asphalt road. Three cost maps were created out of this area (figure 2.8, *top right*): one with a uniform low cost, one with a low cost for the surrounding road, and one with a low cost for the surrounding road and a medium cost for park benches. The other area was grassy with trees, a surrounding outer asphalt road, and a dirt trail cutting straight across. A single cost map was generated from this area, consisting of a low cost for the surrounding road, a high cost for the trees, and a medium cost for the dirt road. Using these different maps, researchers generated a path to navigate between a set of starting and end points with the spiking wavefront algorithm.




Figure 2.7

Path planning using an address event representation table. *Left*: Spike types and neuron IDs are recorded in this table. In order to plan a path using the trained grid of neurons, the neuron corresponding with the location of the agent receives an impulse spike. This spike triggers a wavefront signal to propagate across the grid surface. Since some neurons have longer axonal delays, the wavefront edge travels at different speeds. Using the table, the neuron corresponding to the goal is identified. Then, stepping back through the time steps, a path of neurons can be traced back to the start neuron (*right*). Since costs are encoded using axonal delays, the planned path avoids costlier terrains with obstacles. *Source*: Adapted with permission from Hwu et al. 2018.

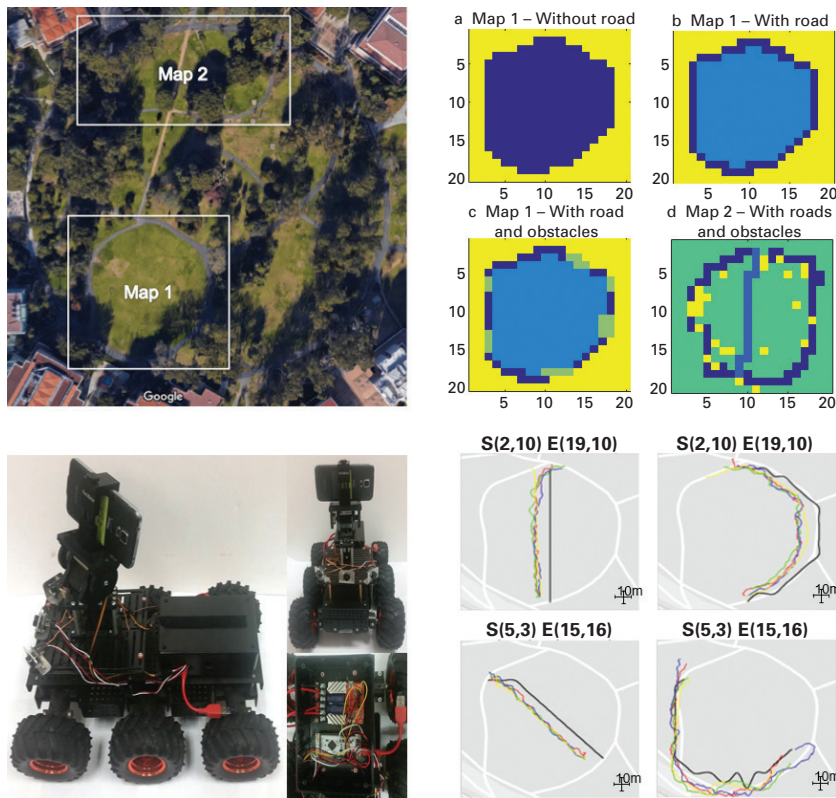


Figure 2.8

Outdoor demonstration of spiking wavefront propagation. *Top left:* A satellite image of an outdoor park where two areas were used to generate cost maps. *Top right:* (a) A uniform cost was given to the grassy area. (b) A low cost was given to the road surrounding the grassy area. (c) A low cost was given to the road, and a medium cost was given to park benches. (d) A low cost was given to the surrounding road, a high cost was given to trees, and a medium cost was given to the dirt path cutting across the area. *Bottom left:* Side, front, and interior views of the Android-based robotics platform. *Bottom Right:* The first row shows two paths planned with the same starting and ending points. The path on the left column was generated using a cost map without the outer road, and the path on the right column was generated using a cost map including the outer road. The bottom row shows the same but with a different set of starting and ending points. When the road is accounted for, the planned path takes the longer, smoother path, as opposed to the shortest path. *Source:* Adapted with permission from Hwu et al. 2018.

Waypoints along the path corresponded to neuronal units representing locations on the map. The robot then used the GPS of the Android phone to drive along the waypoints generated by the algorithm. The paths taken by the robot highlighted trade-offs between finding the shortest path and finding the smoothest path (figure 2.8, *bottom right*). When a uniform cost was used, the shortest path was always chosen. When the road was considered, the robot would occasionally take it, even if it meant traveling a longer distance. For the map containing the dirt road, the robot judged the trade-offs of taking the fastest route versus traveling over bumpy grass. The robot demonstration applied spiking wavefront propagation to cost-aware path planning, showing the possibility of energy savings on an energy-limited mobile platform.

This demonstration combined with the spiking CNN shows the potential for a complete neuromorphic computing solution to outdoor navigation (Hwu, Krichmar, and Zou 2017).

Such a system could enable more computation on mobile platforms and provide more insight into how the brain is able to function with limited energy.

2.5 Future Outlook

Built on a variety of interdisciplinary ideas, neurorobotics has grown into a rich and interesting field. Some of the subtopics of research have remained the same throughout its history, such as navigation, motor planning, mapping, and the development of neural networks. However, the research continues to develop as newer techniques in neurobiology, such as optogenetics, as well as techniques in machine learning and deep neural networks continue to add new tools and insights.

The fields of AI, machine learning, and especially artificial neural networks have enjoyed particular success in recent years. Although deep neural networks have largely been successful, there are a number of new challenges within the field. For the most part, the neural networks work well on specific tasks but have trouble extending knowledge from previously learned tasks to newer but related tasks. Moreover, the neural networks take a large amount of data and training and fail to capture many behaviors that are easy for humans (Larson 2017). This indicates that the study of the brain can contribute much to the field.

According to neuroscientist and entrepreneur Jeff Hawkins (2017), the brain has three key features required for intelligence: 1) learning by rewiring; learning in the brain is both rapid and gradual and can store representations that last over a lifetime; 2) sparse representations; under the constraints of nature, the brain stores information using the fewest metabolic resources possible; 3) embodiment; interaction between the brain and environment together is required for intelligence. We would also argue that the following features are important: 4) value systems; good and bad stimuli from the environment must be learned by detecting saliency and reacting appropriately (Friston et al. 1994; Krichmar 2008) and 5) prediction; we must be able to extrapolate from past experiences to learn how to process future experiences (Clark 2013). Applying these principles, future research in neurorobotics can potentially achieve a more holistic understanding of intelligence, striving for behavior that generalizes across multiple domains and maintains information over long time frames. Neurorobotics is a promising approach to addressing many of the issues the AI community faces today.

2.6 Conclusion

To truly understand intelligence, we believe one must study the brain and body and apply these principles to all applications. Intelligent biological systems are currently our best standard, serving as a model for what AI eventually hopes to achieve. The insights gathered from neurorobotics will ultimately lead to a strong understanding of the essence of intelligence, which will then benefit our understanding of ourselves and lead to applications that improve future technologies.

Additional Reading and Resources

- Edelman, G. M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books. This book introduces an important brain theory that was amenable to testing with neurorobotics.
- Krichmar, J. L., and H. Wagatsuma, eds. 2011. *Neuromorphic and Brain-Based Robots*. Cambridge: Cambridge University Press. This book provides a snapshot of the state of the art in neurorobotics at that time. It covers a range of topics from low-level perception to machine consciousness.
- Tani, Jun. 2016. *Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. Oxford: Oxford University Press. Jun Tani has been a pioneer in neurorobotics. His book covers how higher-order cognition might be realized in neurobots.
- Neurorobotics software and designs:
 - RatSLAM: <https://openslam-org.github.io/openratslam.html>.
 - Android-based robotics platform: <https://www.socsci.uci.edu/~jkrichma/ABR/index.html>.

References

- Almassy, Nikolaus, Gerald M. Edelman, and Olaf Sporns. 1998. "Behavioral Constraints in the Development of Neuronal Properties: A Cortical Model Embedded in a Real-World Device." *Cerebral Cortex* 8 (4): 346–361. <https://doi.org/10.1093/cercor/8.4.346>.
- Arleo, Angelo, and Wulfram Gerstner. 2000. "Modeling Rodent Head-Direction Cells and Place Cells for Spatial Learning in Bio-mimetic Robotics." In Meyer et al. 2000, 236–245.
- Arleo, Angelo, Fabrizio Smeraldi, and Wulfram Gerstner. 2004. "Cognitive Navigation Based on Nonuniform Gabor Space Sampling, Unsupervised Growing Networks, and Reinforcement Learning." *IEEE Transactions on Neural Networks* 15 (3): 639–652. <https://doi.org/10.1109/TNN.2004.826221>. <https://www.ncbi.nlm.nih.gov/pubmed/15384552>.
- Arsenio, Artur M. 2000. "Neural Oscillator Networks for Rhythmic Control of Animats." In Meyer et al. 2000, 105–114.
- Banquet, Jean-Paul, P. H. Gaussier, Mathias Quoy, Arnaud Revel, and Yves Burnod. 2005. "A Hierarchy of Associations in Hippocampo-Cortical Systems: Cognitive Maps and Navigation Strategies." *Neural Computation* 17 (6): 1339–1384. <https://doi.org/10.1162/0899766053630369>. <https://www.ncbi.nlm.nih.gov/pubmed/15901401>.
- Bienenstock, Elie L., Leon N. Cooper, and Paul W. Munro. 1982. "Theory for the Development of Neuron Selectivity: Orientation Specificity and Binocular Interaction in Visual Cortex." *Journal of Neuroscience* 2 (1): 32–48.
- Braitenberg, Valentino. 1986. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
- Bray, Laurence C. Jayet, Sridhar R. Anumandla, Corey M. Thibeault, Roger V. Hoang, Philip H. Goodman, Sergiu M. Dascalu, Bobby D. Bryant, and Frederick C. Harris Jr. 2012. "Real-Time Human-Robot Interaction Underlying Neurobotic Trust and Intent Recognition." *Neural Networks* 32:130–137. <https://doi.org/10.1016/j.neunet.2012.02.029>.
- Browne, William, Kazuhiko Kawamura, Jeffrey Krichmar, William Harwin, and Hiroaki Wagatsuma. 2009. "Cognitive Robotics: New Insights into Robot and Human Intelligence by Reverse Engineering Brain Functions [from the Guest Editors]." *IEEE Robotics and Automation Magazine* 16 (3): 17–18.
- Chavarriaga, Ricardo, Thomas Strössl, Denis Sheynikhovich, and Wulfram Gerstner. 2005. "A Computational Model of Parallel Navigation Systems in Rodents." *Neuroinformatics* 3 (3): 223–241. <https://doi.org/10.1385/NI:3:3:223>. <https://www.ncbi.nlm.nih.gov/pubmed/16077160>.
- Clark, Andy. 2013. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences* 36 (3): 181–204. <https://doi.org/10.1017/S0140525X12000477>. <https://www.ncbi.nlm.nih.gov/pubmed/23663408>.

- Collins, David, and Gordon Wyeth. 2000. "Utilising a Cerebellar Model for Mobile Robot Control in a Delayed Sensory Environment." In *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, edited by Jean-Arcady Meyer, Alain Berthoz, Dario Floreano, Herbert L. Roitblat, and Stewart W. Wilson. Cambridge, MA: MIT Press.
- Davies, Mike, Narayan Srinivasa, Tsung-Han Lin, Gautham China, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou et al. 2018. "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning." *IEEE Micro* 38 (1): 82–99. <https://doi.org/10.1109/MM.2018.112130359>.
- Dragoi, George, and Susumu Tonegawa. 2011. "Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies." *Nature* 469 (7330): 397–401. <https://doi.org/10.1038/nature09633>. <https://www.ncbi.nlm.nih.gov/pubmed/21179088>.
- Edelman, Gerald M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.
- Edelman, Gerald M. 1993. "Neural Darwinism: Selection and Reentrant Signaling in Higher Brain Function." *Neuron* 10 (2): 115–125.
- Edelman, Gerald M., George N. Reeke, W. Einar Gall, Giulio Tononi, Douglas Williams, and Olaf Sporns. 1992. "Synthetic Neural Modeling Applied to a Real-World Artifact." *Proceedings of the National Academy of Sciences* 89 (15): 7267–7271. <https://doi.org/10.1073/pnas.89.15.7267>.
- Esser, Steven K., Paul A. Merolla, John V. Arthur, Andrew S. Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J. Berg et al. 2016. "Convolutional Networks for Fast, Energy-Efficient Neuromorphic Computing." *Proceedings of the National Academy of Sciences* 113 (41): 11441–11446. <https://doi.org/10.1073/pnas.1604850113>. <https://www.ncbi.nlm.nih.gov/pubmed/27651489>.
- Falotico, Egidio, Lorenzo Vannucci, Alessandro Ambrosano, Ugo Albanese, Stefan Ulbrich, Juan Camilo Vasquez Tieck, Georg Hinkel et al. 2017. "Connecting Artificial Brains to Robots in a Comprehensive Simulation Framework: The Neurorobotics Platform." *Frontiers in Neurobotics* 11:2. <https://doi.org/10.3389/fnbot.2017.00002>.
- Fields, R. Douglas. 2015. "A New Mechanism of Nervous System Plasticity: Activity-Dependent Myelination." *Nature Reviews Neuroscience* 16 (12): 756–767. <https://doi.org/10.1038/nrn4023>. <https://www.ncbi.nlm.nih.gov/pubmed/26585800>.
- Floreano, Dario, and Laurent Keller. 2010. "Evolution of Adaptive Behavior in Robots by Means of Darwinian Selection." *PLoS Biol* 8 (1): e1000292. <https://doi.org/10.1371/journal.pbio.1000292>.
- Foster, David J., Richard G. M. Morris, and Peter Dayan. 2000. "A Model of Hippocampally Dependent Navigation, Using the Temporal Difference Learning Rule." *Hippocampus* 10 (1): 1–16. [10.1002/\(Sici\)1098-1063\(2000\)10:1](https://doi.org/10.1002/(Sici)1098-1063(2000)10:1).
- Friston, K. J., Giulio Tononi, G. N. Reeke Jr., Olaf Sporns, and Gerald M. Edelman. 1994. "Value-Dependent Selection in the Brain: Simulation in a Synthetic Neural Model." *Neuroscience* 59 (2): 229–243. <https://www.ncbi.nlm.nih.gov/pubmed/8008189>.
- Gonzalez, F. Montes, Tony J. Prescott, Kevin Gurney, Mark Humphries, and Peter Redgrave. 2000. "An Embodied Model of Action Selection Mechanisms in the Vertebrate Brain." In Meyer et al. 2000, 157–166.
- Hawkins, J. 2017. "What Intelligent Machines Need to Learn from the Neocortex." *IEEE Spectrum* 54 (6): 35–40.
- Hwu, Tiffany, Jacob Isbell, Nicolas Oros, and Jeffrey Krichmar. 2017. "A Self-Driving Robot Using Deep Convolutional Neural Networks on Neuromorphic Hardware." In *2017 International Joint Conference on Neural Networks*, 635–641. New York: IEEE.
- Hwu, Tiffany, Jeffrey Krichmar, and Xinyun Zou. 2017. "A Complete Neuromorphic Solution to Outdoor Navigation and Path Planning." In *2017 IEEE International Symposium on Circuits and Systems*, 1–4. New York: IEEE.
- Hwu, Tiffany, Alexander Y. Wang, Nicolas Oros, and Jeffrey L. Krichmar. 2018. "Adaptive Robot Path Planning Using a Spiking Neuron Algorithm with Axonal Delays." *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 126–137. <https://doi.org/10.1109/TCDS.2017.2655539>. <https://ieeexplore.ieee.org/document/7827038/>.
- Ijspeert, Auke Jan, Alessandro Crespi, and Jean-Marie Cabelguen. 2005. "Simulation and Robotics Studies of Salamander Locomotion." *Neuroinformatics* 3 (3): 171–195. <https://doi.org/10.1385/NI:3:3:171>.
- Ijspeert, Auke Jan, Alessandro Crespi, Dimitri Ryczko, and Jean-Marie Cabelguen. 2007. "From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model." *Science* 315 (5817): 1416–1420. <https://doi.org/10.1126/science.1138353>. <https://www.ncbi.nlm.nih.gov/pubmed/17347441>.
- Indiveri, Giacomo, Bernabé Linares-Barranco, Tara Julia Hamilton, André Van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih-Chii Liu et al. 2011. "Neuromorphic Silicon Neuron Circuits." *Frontiers in Neuroscience* 5:73. <https://doi.org/10.3389/fnins.2011.00073>. <https://www.ncbi.nlm.nih.gov/pubmed/21747754>.
- Kohlbrecher, Stefan, Oskar Von Stryk, Johannes Meyer, and Uwe Klingauf. 2011. "A Flexible and Scalable Slam System with Full 3D Motion Estimation." In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics*, 155–160. New York: IEEE.

- Krichmar, Jeffrey L. 2008. "The Neuromodulatory System: A Framework for Survival and Adaptive Behavior in a Challenging World." *Adaptive Behavior* 16 (6): 385–399.
- Krichmar, Jeffrey L., and Gerald M. Edelman. 2002. "Machine Psychology: Autonomous Behavior, Perceptual Categorization and Conditioning in a Brain-Based Device." *Cerebral Cortex* 12 (8): 818–830. <https://doi.org/10.1093/cercor/12.8.818>.
- Krichmar, Jeffrey L., Douglas A. Nitz, Joseph A. Gally, and Gerald M. Edelman. 2005. "Characterizing Functional Hippocampal Pathways in a Brain-Based Device as It Solves a Spatial Memory Task." *Proceedings of the National Academy of Sciences* 102 (6): 2111–2116. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15684078.
- Krichmar, Jeffrey L., Anil K. Seth, Douglas A. Nitz, Jason G. Fleischer, and Gerald M. Edelman. 2005. "Spatial Navigation and Causal Analysis in a Brain-Based Device Modeling Cortical-Hippocampal Interactions." *Neuroinformatics* 3 (3): 197–221. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=16077159.
- Krichmar, Jeffrey L., James A. Snook, Gerald M. Edelman, and Olaf Sporns. 2000. "Experience-Dependent Perceptual Categorization in a Behaving Real-World Device." In Meyer et al. 2000, 41–50.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. "Imagenet Classification with Deep Convolutional Neural Networks." *Advances in Neural Information Processing Systems* 25:1097–1105. <https://doi.org/10.1145/3065386>.
- Larson, Erik J. 2017. "The Limits of Modern AI: A Story." <https://thebestschools.org/magazine/limits-of-modern-ai/>.
- LaValle, S. M. 2011a. "Motion Planning Part I: The Essentials." *IEEE Robotics and Automation Magazine* 18 (1): 79–89. <https://doi.org/10.1109/Mra.2010.940155>.
- LaValle, S. M. 2011b. "Motion Planning Part II: Wild Frontiers." *IEEE Robotics and Automation Magazine* 18 (2): 108–118. <https://doi.org/10.1109/Mra.2011.941635>.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–444. <https://doi.org/10.1038/nature14539>.
- Liu, Shih-Chii, and Tobi Delbruck. 2010. "Neuromorphic Sensory Systems." *Current Opinion in Neurobiology* 20 (3): 288–295. <https://doi.org/10.1016/j.conb.2010.03.007>. <https://www.ncbi.nlm.nih.gov/pubmed/20493680>.
- Lungarella, Max, Teresa Pegors, Daniel Bulwinkle, and Olaf Sporns. 2005. "Methods for Quantifying the Informational Structure of Sensory and Motor Data." *Neuroinformatics* 3 (3): 243–262. <https://doi.org/10.1385/NI:3:3:243>. <https://www.ncbi.nlm.nih.gov/pubmed/16077161>.
- Mead, Carver. 1990. "Neuromorphic Electronic Systems." *Proceedings of the IEEE* 78 (10): 1629–1636. [10.1109/5.58356](https://doi.org/10.1109/5.58356).
- Merolla, Paul A., John V. Arthur, Rodrigo Alvarez-Icaza, Andrew S. Cassidy, Jun Sawada, Filipp Akopyan, Bryan L. Jackson et al. 2014. "A Million Spiking-Neuron Integrated Circuit with a Scalable Communication Network and Interface." *Science* 345 (6197): 668–673. <https://doi.org/10.1126/science.1254642>. <https://www.ncbi.nlm.nih.gov/pubmed/25104385>.
- Meyer, Jean-Arcady, Alain Berthoz, Dario Floreano, Herbert L. Roitblat, and Stewart W. Wilson, eds. 2000. *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.
- Milford, Michael, Adam Jacobson, Zetao Chen, and Gordon Wyeth. 2016. "RatSLAM: Using Models of Rodent Hippocampus for Robot Navigation and Beyond." In *Robotics Research: The 16th International Symposium ISRR*, 467–485. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-28872-7_27.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. 2015. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518 (7540): 529–533. <https://doi.org/10.1038/nature14236>.
- Mur-Artal, Raul, Jose Maria Martinez Montiel, and Juan D. Tardos. 2015. "ORB-SLAM: A Versatile and Accurate Monocular SLAM System." *IEEE Transactions on Robotics* 31 (5): 1147–1163. <https://doi.org/10.1109/Tro.2015.2463671>.
- Nolfi, Stefano, and D. Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
- O'Keefe, John, and Jonathan Dostrovsky. 1971. "The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat." *Brain Research*. [https://doi.org/10.1016/0006-8993\(71\)90358-1](https://doi.org/10.1016/0006-8993(71)90358-1). <https://www.ncbi.nlm.nih.gov/pubmed/5124915>.
- Pearson, Martin J., Ben Mitchinson, J. Charles Sullivan, Anthony G. Pipe, and Tony J. Prescott. 2011. "Biomimetic Vibrissal Sensing for Robots." *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1581): 3085–3096. <https://doi.org/10.1098/rstb.2011.0164>.

- Pfeifer, R., and J. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Pfeiffer, Brad E., and David J. Foster. 2013. "Hippocampal Place-Cell Sequences Depict Future Paths to Remembered Goals." *Nature* 497 (7447): 74–79. <https://doi.org/10.1038/nature12112>. <https://www.ncbi.nlm.nih.gov/pubmed/23594744>.
- Reeke, George N., Olaf Sporns, and Gerald M. Edelman. 1990. "Synthetic Neural Modeling: The 'Darwin' Series of Recognition Automata." *Proceedings of the IEEE* 78 (9): 1498–1530. <https://doi.org/10.1109/5.58327>.
- Schultz, Wolfram, Peter Dayan, and P. Read Montague. 1997. "A Neural Substrate of Prediction and Reward." *Science* 275 (5306): 1593–1599. <https://doi.org/10.1126/science.275.5306.1593>. <https://www.ncbi.nlm.nih.gov/pubmed/9054347>.
- Seth, Anil K., Jeffrey L. McKinstry, Gerald M. Edelman, and Jeffrey L. Krichmar. 2004a. "Texture Discrimination by an Autonomous Mobile Brain-Based Device with Whiskers." In Vol. 5, *Proceedings of the IEEE International Conference on Robotics and Automation*, 4925–4930. New York: IEEE.
- Seth, Anil K., Jeffrey L. McKinstry, Gerald M. Edelman, and Jeffrey L. Krichmar. 2004b. "Visual Binding through Reentrant Connectivity and Dynamic Synchronization in a Brain-Based Device." *Cerebral Cortex* 14 (11): 1185–1199.
- Seth, Anil K., Olaf Sporns, and Jeffrey L. Krichmar. 2005. "Neurorobotic Models in Neuroscience and Neuroinformatics." *Neuroinformatics* 3:167–170.
- Stewart, Terrence C., Ashley Kleinhans, Andrew Mundy, and Jörg Conradt. 2016. "Serendipitous Offline Learning in a Neuromorphic Robot." *Frontiers in Neurobotics* 10:1. S0954–898x(02)36091–3.
- Stringer, S. M., T. P. Trappenberg, E. T. Rolls, and I. E. de Araujo. 2002. "Self-Organizing Continuous Attractor Networks and Path Integration: One-Dimensional Models of Head Direction Cells." *Network* 13 (2): 217–242. PMID: 12061421.
- Zhang, Honghui, and Joshua Jacobs. 2015. "Traveling Theta Waves in the Human Hippocampus." *Journal of Neuroscience* 35 (36): 12477–12487. <https://doi.org/10.1523/JNEUROSCI.5102-14.2015>. <https://www.ncbi.nlm.nih.gov/pubmed/26354915>.
- Zilli, Eric A. 2012. "Models of Grid Cell Spatial Firing Published 2005–2011." *Frontiers in Neural Circuits* 6:16. <https://doi.org/ARTN>.

3 Developmental Robotics

Minoru Asada and Angelo Cangelosi

3.1 Introduction

In this chapter we introduce “developmental robotics” in the context of cognitive robotics. Developmental robotics can be defined as “the interdisciplinary approach to the autonomous design of behavioral and cognitive capabilities in artificial agents (robots) that takes direct inspiration from the developmental principles and mechanisms observed in the natural cognitive systems of children” (Cangelosi and Schlesinger 2015, 4). Developmental robotics relies on a highly interdisciplinary effort of developmental psychology, neuroscience, and comparative psychology with robotics and artificial intelligence. In particular, developmental sciences such as child psychology provide the empirical bases to identify the general developmental principles, mechanisms, models, and experimental protocols guiding the design of cognitive robots and their testing in situated developmental robotics experiments. Given this close interaction, developmental psychology and developmental robotics can also mutually benefit from such a combined effort (Cangelosi and Schlesinger 2018).

Developmental robotics is based on the vision that a baby robot, using developmental principles and mechanisms regulating the real-time interaction between its body, brain, and environment, can autonomously acquire an increasingly complex set of sensorimotor and mental capabilities. Thus, within the wider approach of cognitive robotics, developmental robotics specializes in its emphasis on the design of baby robots with an autonomous capability to acquire ever-more-complex skills.

Historically, the field of developmental robotics has also been known as “cognitive developmental robotics” (Asada et al. 2001), “autonomous mental development” (Weng et al. 2001), and “epigenetic robotics” (Zlatev and Balkenius 2001). Asada et al. (2001) proposed “cognitive developmental robotics” as a new paradigm for the design of humanoid robots. Lungarella et al. (2003) published the first survey paper on developmental robotics. Asada et al. (2009) later proposed a systematic survey of the early cognitive developmental robotics approaches. More recently, Cangelosi and Schlesinger (2015) provided a comprehensive review of the field in their book *Developmental Robotics: From Babies to Robots*.

In this chapter we will first consider the theoretical background of cognitive developmental robotics, focusing on epistemological paradigm shifts from human-object dichotomy

to human-machine physical and mental interaction. Based on this background, “physical embodiment” and “social interaction” are introduced as key concepts of developmental robotics (Asada et al. 2009). We will then extend this to the six defining principles of developmental robotics proposed by Cangelosi and Schlesinger (2015), with brief examples of each.

3.2 Theoretical and Philosophical Background

Asada (2019) has proposed a general outline of the theoretical and philosophical background of the relationship between consciousness, humans, and objects/technology at the origin of cognitive developmental robotics. The discussion below follows the concepts introduced in Tani (2016) but adds further consideration of the contribution of the philosophers Kant and Vico (figure 3.1).

Initially, Descartes advanced mind-body dualism,¹ establishing the relationship between mind and body or things, and laid the foundation for modern philosophy. Then Vico opposed Cartesianism and all reductionism, asserting the *verum factum* principle that truth is verified only by creation or invention, not by observations, as proposed in Cartesianism.²

Husserl, Heidegger, and Merlot-Ponty presented important concepts such as embodiment, interaction, and intersubjectivity (see Tani 2016) and noted that the essence of reality is lost by discriminating between humans and objects (cf. Asada [2019] for more details on this issue).

In his moral philosophy, Kant spoke from the perspective of morality as an obligatory act—that is, “what should be.”³ In today’s world, due to technological progress, interacting with objects exposes the limits of anthropocentric thinking. Peter-Paul Verbeek has shown a typical example of such a situation when he and his wife entered the ultrasound examination room. He mentioned in the preface of *Moralizing Technology: Understanding and Designing the Morality of Things* that “even though the technology in the ultrasound practice clearly had moral significance, it did not directly steer our behavior. Rather, it helped to shape our experience of our unborn child and the interpretive frameworks that guided our actions and decisions. By us, this technology had not simply granted us a ‘peek into the womb’; it had reorganized the relations between our unborn child and ourselves” (Verbeek 2011). As he mentioned, this is one aspect of the moral significance of technology, and the fixed view of Kant’s moral philosophy seems unable to handle appropriate relationships with these technologies. Foucault’s (1994) moral ethics, or “what we like to be,” are considered more relevant.

Figure 3.1 shows a paradigm shift from mind-body dualism, which emphasizes the relationship between humans and objects by anthropocentric thinking (above the thick broken line in the figure), to a concept that emphasizes both the importance of a creation-based viewpoint and societal impacts (below the thick broken line in the figure). In other words, objects and technologies have come to judge and commit to decision-making via machine learning represented by deep learning, such as autodrive, and the structure and mechanism of free will and consciousness have gradually been revealed in neuroscience, physiology, and cognitive sciences. This is because traditional views of consciousness and autonomy no longer function in modern disciplines.




Figure 3.1
 Outline of the philosophical background of the relationship between humans and things (technology).

All these theoretical considerations have significantly influenced the approach of cognitive developmental robotics (Asada et al. 2009; Asada 2019). These epistemological concepts advocate the importance of physical embodiment and social interaction, which have influenced the wider field of cognitive robotics, as discussed in chapter 1. Before introducing the key principles and related studies of developmental robotics, we review the developmental process of the human fetus and infant, which will have an impact on the design issues and approaches of developmental robotics.

3.3 A Brief Overview of the Development of the Human Fetus and Infant

Advanced imaging technologies such as three-dimensional ultrasound movies have enabled the observation of various kinds of fetal movements in the womb after several weeks of gestation. This reveals the possibility of the fetus learning in the womb (Hopson 1998). De Vries et al. (1984) reported that fetal motility started from the early state of “just discern movements (7.5 weeks)” to the later state of “sucking and swallow (12.5–14.5 weeks)” through “startle, general movements, hiccup, isolated arm movements, isolated leg movements, head retroflexion, head rotation, hand/face contact, breathing movements, jaw opening, stretch, head anteflexion, and yawn.” Campbell (2004) also reported that the eyes of the fetus open around twenty-six weeks’ gestation and that the fetus often touches their face with the hands during embryonic weeks twenty-four and twenty-seven.

Touch is the first sense to develop in the fetus, followed by the other senses, such as taste, hearing, and vision. Chamberlain stated that just before eight weeks’ gestational age, the first sensitivity to touch manifests in a set of protective movements to avoid a mere hair stroke on the cheek. From this early stage, experiments with a hair stroke on various parts of the body show that skin sensitivity quickly extends to the genital area (ten weeks), palms (eleven weeks), and soles (twelve weeks). These areas of first sensitivity will have the greatest number and variety of sensory receptors in the adult. By seventeen weeks, all parts of the abdomen and buttocks become sensitive. Skin is marvelously complex, containing a hundred varieties of cells that seem especially sensitive to heat, cold, pressure, and pain. By thirty-two weeks, nearly every part of the body is sensitive to the same light stroke of a single hair. Both hearing and vision start to function about eighteen weeks after gestation and fully develop at around twenty-five weeks.

Moreover, it is reported that visual stimulation from the outside of the maternal body can activate the fetal brain (Eswaran et al. 2002). Figure 3.2 summarizes the emergence of fetal movements with the development of the fetal senses reviewed above.

After birth, infants are supposed to gradually develop body representation, categories for graspable objects, capability of mental simulation of actions, and so on through learning processes. For example, controlling the hand at the fifth month means learning the forward and inverse models of the hand. Table 3.1 shows typical behaviors and their corresponding targets to learn.

Our growing understanding of the early stages of fetus and infant development have been very influential in developmental robotics. Asada et al. (2009) analyzed in detail a wide set of pioneering developmental robotics models of early fetal and infant development. Two three-dimensional simulation models of the fetus and newborn infants were developed by




Figure 3.2
Emergence of fetal movements and the senses. *Source:* Adapted from fig. 1 in de Vries et al. 1984.

Table 3.1
Infant developmental behavior and learning targets

Month—behavior	Learning targets
5 hand regard	Forward and inverse models of the hand
6 finger the another's face	Integration of visuotactile sensation of the face
7 drop objects and observe the result	Causality and permanency of objects
8 hit objects	Dynamic modeling of objects
9 drum or bring a cup to mouth	Tool use
10 imitate movements	Imitation of unseen movements
10 rudimentary sympathy	Feel pain and empathy
11 grasp and carry objects to others	Action recognition and generation, cooperation
12 pretend	Mental simulation

Source: Adapted from Asada et al. 2009.

Kuniyoshi and colleagues within the Japan Science and Technology Agency Exploratory Research for Advanced Technology (JST ERATO) Asada project. The first model (Kuniyoshi and Sangawa 2006) provided the initial, minimally simple body model of fetal and neonatal development. The subsequent fetus model (Mori and Kuniyoshi 2010) produced a more realistic rendering of the fetus's sensorimotor apparatus and a stronger focus on learning experiments. These models offer a useful research tool to investigate prebirth sensorimotor development by providing a realistic representation of the fetus's sensors and the reaction of the body to gravity and the womb border and environment. The first model, for instance, was used to study the role of general embodied developmental principles in early sensorimotor learning. In particular, it aimed at exploring the hypothesis that partially ordered embodiment dynamical patterns emerge from the chaotic exploration of body-brain-environment interactions during gestation. These patterns lead, later in development, to the emergence of meaningful motor behavior such as rolling over and crawling motions in neonates.

3.4 Six Principles of Developmental Robotics

Figure 3.3 shows the six principles of developmental robotics, centering two key concepts, embodiment and social interaction. Embodiment, or physical embodiment, is a fundamental constraint for infants (humans and robots) to learn sensorimotor mapping through interaction with the environment. Related research topics are motor babbling and body representation (body schema or body image) through crossmodal association (e.g., Mannella et al. 2018). These topics lead to the emergence of the early concept of the self, often called the “ecological self” (Neisser 1994) through embodied, situated, and enactive development. The ecological self is also called the temporary self or, according to Gallagher (2000), the minimal self, involving a sense of agency or a sense of ownership of motion.

The early stage of social interaction can be observed as infant-caregiver interaction. Intrinsic motivation and social-learning instinct (Baldassarre and Mirolli 2013; Ishihara et al. 2011) inside the agents play important roles in developing various behavioral and cognitive functions, such as imitation (e.g., gesture, vocalization, and joint attention), turn taking, and so on.

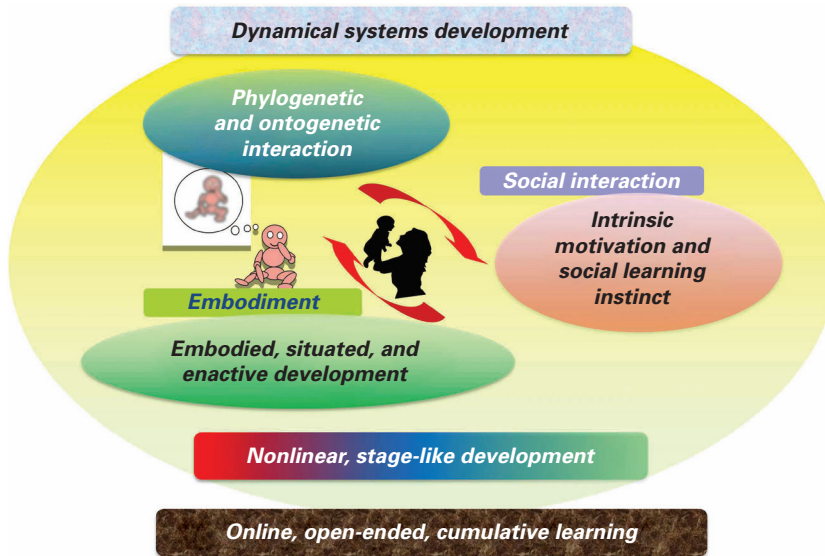


Figure 3.3
Principles of developmental robotics.

Both phylogenetic and ontogenetic interactions occur during the above developmental processes. Innate functions are regarded as assumptions, and learning targets are set at each stage of development. The learning results become the assumptions for the next stage of learning, and vice versa—that is, the assumptions at the current stage might be the results of learning during the previous stage. Thus, nonlinear stagelike learning develops (Lee et al. 2007).

Developmental pathways are diverse, from typical development to atypical, and this also holds true for developmental robots. These pathways are expected to share several key points that enable social interactions from different pathways, and learning continues beyond different stages in terms of a lifelong scale as a whole. It is an online, open-ended, cumulative learning process.

Thelen and Smith (1994) proposed the dynamical systems approach as a developmental psychology theory, and several computational-modeling methods attempt to reproduce nonlinear, dynamic developmental processes of coupled interactions involving the classical nature-nurture issue.

In the following sections, we will describe in detail the six key defining principles of developmental robotics, as proposed in Cangelosi and Schlesinger (2015). The presentation of each principle will refer to certain seminal developmental psychology studies and related developmental robotics models.

3.4.1 Dynamical Systems Development

In mathematics, a dynamical system is characterized by complex changes, over time, in the phase state that result from the self-organization of multifaceted interactions between the system's variables. The complex interaction of nonlinear phenomena results in the production of unpredictable states of the system, often referred to as emergent states. In child psychology this concept has been borrowed by Thelen and Smith (1994) to explain child

development as the emergent product of the intricate and dynamic interaction of many decentralized and local interactions related to the child's growing body and brain and the environment. Thus, Thelen and Smith have proposed that the development of a child should be viewed as change within a complex dynamic system, where the growing child can generate novel behavior through interaction with the environment, and these behavioral states vary in their stability within the complex system.

One key concept in this theory is that of *multicausality*—for example, in the case when one behavior, such as crawling and walking, is determined by the simultaneous and dynamic consequences of various phenomena at the level of the brain, body, and environment. Thelen and Smith analyzed the dynamic changes in crawling and walking as an example of multicausality changes in the child's adaptation to the environment, in response to body growth. When the child's body configuration produces sufficient strength and coordination to support them through the hands and knees posture but is not strong enough for upright walking, the child settles for a crawling strategy to locomote in the environment. But when the infant's body growth results in stronger and more stable legs, the standing and walking behavior emerges as the stable developmental state, which as a consequence destabilizes, and gradually stops, the pattern of crawling. This demonstrates that the locomotion behavior is the result of self-organizing dynamics of decentralized factors such as the child's changing body (stronger legs and better balance) and its adaptation to the environment.

Another key concept in the dynamical systems view of development is that of *nested timescales*. That is, neural and embodiment phenomena act at different timescales and affect development in an intricate, dynamical way. For example, the dynamics of the very fast timescale of neural activity (milliseconds) is nested within the dynamics of the other slower timescales, such as action-reaction time (seconds or hundreds of milliseconds), learning (after hours or days), and physical body growth (months). One of the best-known developmental psychology examples used by Thelen and Smith to demonstrate the combined effects of the concepts of multicausality and nested timescales is that of the A-not-B error. This example is inspired by Piaget's object permanence experiment, when one toy is repeatedly hidden under a lid at a location A (right) during the first part of the experiment and then, toward the end of the task, is hidden in a location B (left) for a single trial, and the child is asked to reach for the object. While infants older than twelve months have no problem in reaching for the toy in its correct location B, unexpectedly, most eight-to-ten-month-old children err in looking for the object in location A. Although psychologists such as Piaget have used explanations based on age (stage) differences linked to qualitative changes in the ability to represent objects and space, a computational simulation of the dynamical system model (Thelen et al. 2002) has demonstrated that many decentralized factors (multicausality) and timing manipulations (nested timing) affect such a situation. These, for example, depend on the time delay between hiding and reaching, the properties of the lids on the table, the saliency of the hiding event, and the past activity of the infant and their body posture.

The use of a dynamical system approach as a theory of development has had significant influence in developmental robotics research, as well as in other cognitive robotics areas (Beer 2000; Nolfi and Floreano 2000). This theory has been applied, for example, to developmental robotics models of early motor development, as in Mori and Kuniyoshi's

(2010) simulation on the self-organization of body representation and general movements in the fetus and newborn. In Meola et al. (2015) and Mannella et al. (2018), the initial dynamical movements of a robot, analogous to Piaget's circular reactions, were progressively shaped into purposeful actions. Additionally, a developmental robotics model of early word learning (Morse et al. 2010) uses a similar setup to the A-not-B error to investigate dynamic interactions between embodiment factors and higher-order language development phenomena. Tani (2016) also showed approaches to neurorobotics based on the idea of dynamical systems.

3.4.2 Embodied, Situated, and Enactive Development

Chapter 1 has already discussed the role of *embodiment* in robot design. In addition, two more concepts have influenced developmental robotics models. One is the role of interaction between the body and its environment (*situatedness*), and the other looks at the organism's autonomous acquisition of a model of the world through sensorimotor interactions (*enaction*).

Ziemke (2001) and Wilson (2002) analyzed different views of embodiment and their consideration in computational models and psychology experiments. These views ranged from considering embodiment as the phenomenon of “structural coupling” between the body and the environment to the more restrictive “organismic” embodiment view based on the autopoiesis of living systems—that is, that cognition actually is what living systems do in interaction with their world (Varela et al. 1992). Along the same lines, the paradigm of enaction highlights the fact that an autonomous cognitive system interacting in its environment is capable of developing its own understanding of the world and generating its own models of how the world works (Vernon 2010; Stewart et al. 2010).

Embodied and situated intelligence has significantly influenced developmental robotics, and practically any developmental model places great emphasis on the relation between the robot's body, brain, and environment. Embodiment effects concern pure motor capabilities (morphological computation) as well as higher-order cognitive skills such as language (grounding) and imagination. Hoffmann et al. (2010) surveyed various approaches to body representations in robotics. Among them, body image/schema acquisition by Yoshikawa et al. (2002), Fuke et al. (2007), and Hikita et al. (2008) focused on crossmodal association and self-organizing maps, both of which are powerful methods in developmental robotics. Yamada et al. (2016) showed a brain-body interaction in the fetus utilizing 2.6 million spike neurons and a realistic musculoskeletal model, although it was computer simulation.

On the other end, an example of the role of embodiment in higher-order cognitive functions can be seen in models of the grounding of words in action and perception (Cangelosi 2010; Morse et al. 2010), the relationship between spatial representation and numerical cognition in psychology and developmental robotics (Rucinski et al. 2011; see also chapter 22), and the relationship between sensorimotor behavior and imagination processes (Seepanomwan et al. 2015).

3.4.3 Intrinsic Motivation and Social-Learning Instinct

Developmental robotics explores methods for designing *intrinsically motivated* agents and robots who can define their own goals and value systems (see chapter 13; Baldassarre and

Mirolli 2013). An intrinsically motivated robot explores its environment in a completely autonomous manner by deciding for itself what it wants to learn and what goals it wants to achieve. In other words, intrinsic motivation enables the agent to construct its own value system.

The concept of intrinsic motivation is inspired by a variety of behaviors and skills that begin to develop in infancy and early childhood, including diverse phenomena such as curiosity, surprise, novelty seeking, and the “drive” to achieve mastery. Oudeyer et al. (2007) proposed a framework for organizing research on models of intrinsic motivation, including two major categories: 1) knowledge-based approaches (later subdivided into novelty-based and prediction-based approaches; Barto et al. 2013) and 2) competence-based approaches. Within this framework, a large number of algorithms can be defined and systematically compared.

Novelty-based approaches to intrinsic motivation study robots that learn about their environments by exploring and discovering unusual or unexpected features. A useful mechanism for detecting novelty is habituation: the robot compares its current sensory state to past experiences, devoting its attention to situations that are unique or different (e.g., Vieira Neto and Nehmzow 2007).

Prediction-based approaches use knowledge-based intrinsic motivation to explicitly attempt to predict future states of the world (Schmidhuber 2010). The rationale of this approach is that incorrect or inaccurate predictions provide a learning signal—that is, they indicate events that are poorly understood and require further analysis and attention. As an example of this approach, Oudeyer et al. (2005) describe the playground experiment, in which the Sony AIBO robot learned to explore and interact with a set of toys in its environment.

The third approach to modeling intrinsic motivation is competence based. The robot is motivated to explore and develop skills that effectively produce reliable consequences (Barto et al. 2004; Santucci et al. 2016). A key element of the competence-based approach is contingency detection (Jacquey et al. 2019): this is the capacity to detect when one’s actions have an effect on the environment. While the knowledge-based approaches motivate the agent toward discovering properties of the world, the competence-based approach, in contrast, motivates the agent to discover what it can do with the world.

Child development research has shown the presence of social-learning capabilities (instincts). This is evidenced, for example, by observations that newborn babies instinctually imitate the behavior of others from the very first day of life and can imitate complex facial expressions (Meltzoff and Moore 1977). Moreover, comparative psychology studies have demonstrated that eighteen-to-twenty-four-month-old children have a tendency to cooperate altruistically, a capacity not observed in chimpanzees (Warneken et al. 2006).

Developmental robotics places a heavy emphasis on social learning; various robotics models of joint attention, imitation, and cooperation have been tested. Nagai et al. (2003b, 2006) showed the early developmental model of joint attention, and Sumioka et al. (2010) proposed a contingency model for joint attention. Asada (2016) reviewed modeling approaches to early vocal development through infant-caregiver interaction. Imitation and cooperation have been other hot topics in general, with representative studies introduced by Cangelosi and Schlesinger (2015).

3.4.4 Phylogenetic and Ontogenetic Interaction

Two different timescales must be considered in developmental robotics: 1) the ontogenetic phenomena of learning, over a timescale of hours or days, with maturational changes occurring for periods of months or years and 2) the phylogenetic phenomena of evolutionary changes. Therefore, the additional implication of the interaction between ontogenetic and phylogenetic phenomena should be considered in developmental robotics models of development.

The whole process of development can be observed as a heterogeneous interaction between phylogenetic constraints and ontogenetic processes. Therefore, the issues are not a simple dichotomy of “nature versus nurture.” Ridley (2003) reframed this dichotomy in terms of “nature via nurture.” Although it has been said that “ontogeny recapitulates phylogeny,” it does not seem so simple. Both processes are highly intertwined and show a broad and dynamic landscape as a result.

Maturation refers to changes in the anatomy and physiology of both the child’s brain and body, especially during the first years of life. Maturation phenomena related to the brain include the decrease of brain plasticity during early development and the gradual hemispheric specialization and pruning of neurons and connections (Abitz et al. 2007). Brain maturation changes have also been evoked to explain the critical period in learning. Critical periods are stages (windows of time) of an organism’s life span during which the individual is more sensitive to external stimulation and more efficient at learning. Moreover, after the critical period has ended, learning becomes difficult or impossible. The best-known example of a critical period (aka a “sensitive period”) in ethology is Konrad Lorenz’s study on imprinting—that is, the attachment of ducklings to their mother (or to Lorenz), which is only possible within the first few hours of life and has a long-lasting effect. In vision research, Hubel and Wiesel (1970) demonstrated that the cat’s visual cortex can only develop its receptive fields if the animal is exposed to visual stimuli during the first few months of life and not when it experiences total visual deprivation as a kitten by having its eyes covered.

Maturation in the body of the child is more evident, given the significant morphological changes a child goes through from birth to adolescence. These changes naturally affect the motor development of the child, as in Thelen and Smith’s (1994) analysis of crawling and walking. Morphological changes occurring during development also have an implication for the exploitation of embodiment factors, as discussed in 3.4.2.

Some developmental robotics models have explicitly addressed the issue of brain and body maturation changes. For example, the study by Schlesinger et al. (2007) modeled the effects of neural plasticity in the development of object perception skills.

Ontogenetic changes due to maturation and learning have important implications for the interaction of development with phylogenetic changes due to evolution. Body morphology and brain plasticity variations can in fact be explained as evolutionary adaptations of the species to changing environmental contexts. These phenomena have been analyzed in terms of genetic changes affecting the timing of ontogenetic phenomena, known as *heterochronic changes* (McKinney et al. 1991). Heterochronic changes have been used to explain the complex interaction between nature and nurture in models of development, as in Elman et al.’s (1996) proposal that the role of genetic factors in development is to determine the

architectural constraints, which subsequently control learning. Such constraints can be explained in terms of brain adaptation and neurodevelopmental and maturational events.

The interaction between ontogenetic and phylogenetic factors has been investigated through computational modeling. For example, Hinton and Nowlan (1987) and Nolfi et al. (1994) have developed simulation models explaining the effects of learning in evolution, as in the Baldwin effect. Cangelosi (1999) tested the effects of heterochronic changes in the evolution of neural network architectures for simulated agents. Furthermore, the modeling of the evolution of varying body and brain morphologies in response to phylogenetic and ontogenetic requirements is also the goal of the “evo-devo” computational approach. This aims to simulate the simultaneous effects of developmental and evolutionary adaptation in body and brain morphologies (e.g., Stanley and Miikkulainen 2003; Kumar and Bentley 2003; Pfeifer and Bongard 2006).

Developmental robotics models are normally based on robots with fixed morphologies and cannot directly address the simultaneous modeling of phylogenetic changes and its interaction with ontogenetic morphological changes. However, various epigenetic robotics models take into consideration the evolutionary origins of the ontogenetic changes of learning and maturation, especially for studies including changes in brain morphology. Nagai et al. (2006) compared performances in terms of the timing parameter that controls the learning phase for joint attention. One is a fixed time schedule in which the learning phase shifts to the next one (phylogenetic constraint), and the other depends on the learning result—that is, shifting to the next learning phase if the target of the current learning phase is achieved (ontogenetic constraint). Because this is a case of brain maturation, the fixed time schedule for the timing parameter could be arbitrary, allowing the designer to tune anyway. However, in the case of body maturation it may interfere with the learning process, accelerating or decelerating the process to some extent. This seems to be a typical case of a heterogeneous interaction between the phylogenetic and ontogenetic processes. Although the brain maturation process was not clearly described, Yamada et al. (2010) showed two kinds of computer simulations for the fetus and the infant that indicated body maturation.

3.4.5 Nonlinear, Stagelike Development

The literature on child psychology has plenty of theories and models proposing a sequence of developmental stages. Each stage is characterized by the acquisition of specific behavioral and mental strategies, which become more complex and articulated as the child progresses through these stages. Piaget’s (1952) four stages of development of thought is the prototypical example of a theory of development centered on stages. Numerous other examples of stage-based development exist (Courage and Howe 2002; Butterworth and Jarrett 1991).

In most theories, the transition between stages follows nonlinear, qualitative shifts. In the example of Piaget’s four stages, the mental schemas used in each stage are qualitatively different, as they are the results of accommodation processes that change and adapt the schema to new knowledge representations and operations. Another well-known developmental theory based on qualitative changes during development is the representational-redescription model of Karmiloff-Smith (1995). Her model assumes four levels of development going from the use of implicit representation to different levels of explicit knowledge-representation strategies. When a child learns new facts and knowledge about specific domains, they develop new representations, which are gradually “redescribed” and increase the child’s explicit

understanding of the world. This has been applied to a variety of knowledge domains such as physics, math, and language.

The nonlinearity of the developmental process, and the qualitative shifts in the mental strategies and knowledge representations employed by the child at different stages of development, has been extensively investigated through U-shaped learning-error patterns and the vocabulary spurt phenomenon—that is, the sudden growth of the vocabulary after a slow word-acquisition stage (Elman et al. 1996).

Many developmental robotics studies aim to model the progression of stages during the robot's development, with some directly addressing the issue of nonlinear phenomena in developmental stages as a result of learning dynamics. Ogino et al. (2006) proposed an active lexicon-acquisition method based on curiosity to partially model the vocabulary spurt phenomenon. Nagai et al. (2003a) explicitly modeled the joint attention stages proposed by Butterworth and Jarrett (1991). However, the model shows that qualitative changes between these stages are the result of gradual changes in the robot's neural and learning architecture, rather than ad hoc manipulations of the robot's attention strategies. Some models have also directly addressed the modeling of U-shaped phenomena, such as the Morse et al. (2011) model of error patterns in phonetic processing. Asada (2015) proposed a conceptual model for the development of artificial empathy that shows a stagelike development starting from emotional contagion through emotional/cognitive empathy to sympathy/compassion. Lee et al. (2007) proposed the lift constraint, act, and saturate method for which robots can develop increasingly complex skills by “saturating” the acquisition of knowledge at a certain level of competence and thus release the possibility of learning at a more complex level.

3.4.6 Online, Open-Ended, Cumulative Learning

Human development is characterized by online, crossmodal, continuous, open-ended learning. “Online” refers to the fact that learning happens while the child interacts with the environment and not in an off-line mode. “Crossmodal” refers to the fact that different modalities and cognitive domains are acquired in parallel by the child and interact with each other. “Continuous” and “open-ended” refers to the fact that learning and development do not start and stop at specific stages but rather are lifelong learning experiences (Baldassarre and Mirolli 2013).

Online learning is currently implemented in developmental robotics. However, the application of crossmodal, cumulative, open-ended learning, which can lead to cognitive bootstrapping phenomena, has been investigated less frequently. Most of the current models typically focus on the acquisition of only one task or modality (perception, or phonetics, or semantics, and so on), and few consider the parallel development, and interaction, between modalities and cognitive functions. Thus, a truly online, crossmodal, cumulative, open-ended developmental robotics model remains a fundamental challenge for the field.

3.5 Conclusion

The numerous philosophical consideration and research issues, challenges, and principles discussed have led to the creation of numerous developmental robotics models exploring a wide range of behavioral and cognitive skills. In many of the chapters of part III, which

focus on cognitive robotics models of specific sensorimotor and cognitive functions, we will see further examples of developmental robotics models and experiments. For example, chapter 13 is largely based on developmental approaches, and chapter 18 and 20 present various developmental robotics models of social and linguistic skills.

Additional Reading and Resources

- The most comprehensive overview of the field of developmental robotics: Cangelosi, Angelo, and Schlesinger, Matthew. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Seminal review paper on the initial theoretical issues and pioneering models of baby robots: Asada, Minoru, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. 2009. “Cognitive Developmental Robotics: A Survey.” *IEEE Transactions on Autonomous Mental Development* 1 (1): 12–34.
- A rich theoretical and computational analysis of principles and models of cognitive and developmental robotics: Tani, Jun. 2016. *Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. Oxford: Oxford University Press.

Notes

1. Substance dualism, material-centered dualism, spiritual dualism, classical dualism, and so on, <https://www.iep.utm.edu/dualism/>.
2. Wikipedia, “Giambattista Vico,” https://en.wikipedia.org/wiki/Giambattista_Vico.
3. Wikipedia, “Kantian Ethics,” https://en.wikipedia.org/wiki/Kantian_ethics.

References

- Abitz, Maja, Rune Damgaard Nielsen, Edward G. Jones, Henning Laursen, Niels Graem, and Bente Pakkenberg. 2007. “Excess of Neurons in the Human Newborn Mediodorsal Thalamus Compared with That of the Adult.” *Cerebral Cortex* 17 (11): 2573–2578.
- Asada, Minoru. 2015. “Towards Artificial Empathy.” *International Journal of Social Robotics* 7:19–33.
- Asada, Minoru. 2016. “Modeling Early Vocal Development through Infant-Caregiver Interaction: A Review.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (2): 128–138.
- Asada, Minoru. 2019. “Artificial Pain May Induce Empathy, Morality, and Ethics in the Conscious Mind of Robots.” *Philosophies* 4:38–47.
- Asada, Minoru, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. 2009. “Cognitive Developmental Robotics: A Survey.” *IEEE Transactions on Autonomous Mental Development* 1 (1): 12–34.
- Asada, Minoru, Karl F. MacDorman, Hiroshi Ishiguro, and Yasuo Kuniyoshi. 2001. “Cognitive Developmental Robotics as a New Paradigm for the Design of Humanoid Robots.” *Robotics and Autonomous Systems* 37:185–193.
- Baldassarre, Gianluca, and Marco Mirolli, eds. 2013. *Intrinsically Motivated Learning in Natural and Artificial Systems*. Berlin: Springer.
- Barto, Andrew, Marco Mirolli, and Gianluca Baldassarre. 2013. “Novelty or Surprise?” *Frontiers in Psychology* 4:907.
- Barto, Andrew G., Satinder Singh, and Nuttapong Chentanez. 2004. “Intrinsically Motivated Learning of Hierarchical Collections of Skills.” In *Proceedings of the 3rd International Conference on Development and Learning*, 112–119. La Jolla, CA: UCSD Institute for Neural Computation.
- Beer, Randall D. 2000. “Dynamical Approaches to Cognitive Science.” *Trends in Cognitive Science* 4 (3): 91–99.

- Butterworth, G. E., and N. L. M. Jarrett. 1991. "What Minds Have in Common Is Space: Spatial Mechanisms Serving Joint Visual Attention in Infancy." *British Journal of Developmental Psychology* 9:55–72.
- Campbell, Stuart. 2004. *Watch Me Grow: A Unique, 3-Dimensional Week-by-Week Look at Your Baby's Behavior and Development in the Womb*. New York: Macmillan.
- Cangelosi, Angelo. 1999. "Heterochrony and Adaptation in Developing Neural Networks." In *Proceedings of GECCO99 Genetic and Evolutionary Computation Conference*, 1241–1248. San Francisco: Morgan Kaufmann Publishers.
- Cangelosi, Angelo. 2010. "Grounding Language in Action and Perception: From Cognitive Agents to Humanoid Robots." *Physics of Life Reviews* 7:139–151.
- Cangelosi, Angelo, and Matthew Schlesinger. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Cangelosi, Angelo, and Matthew Schlesinger. 2018. "From Babies to Robots: The Contribution of Developmental Robotics to Developmental Psychology." *Child Development Perspectives* 12 (3): 183–188.
- Courage, Mary L., and Mark L. Howe. 2002. "From Infant to Child: The Dynamics of Cognitive Change in the Second Year of Life." *Psychological Bulletin* 128:250–277.
- de Vries, J. I. P., G. H. A. Visser, and H. F. R. Prechtl. 1984. "Fetal Motility in the First Half of Pregnancy." *Clinics in Developmental Medicine* 94:46–64.
- Elman, Jeffrey L., Elizabeth A. Bates, Mark H. Johnson, Annette Karmiloff-Smith, Kim Plunkett, and Domenico Parisi. 1996. *Rethinking Innateness: A Connectionist Perspective on Development*. Vol. 10. Cambridge, MA: MIT Press.
- Eswaran, Hari, James D. Wilson, Hubert Preissl, Stephen E. Robinson, Jiri Vrba, Pam Murphy, Douglas F. Rose, and Curtis L. Lowery. 2002. "Magnetoencephalographic Recordings of Visual Evoked Brain Activity in the Human Fetus." *Lancet* 360 (9335): 779–780.
- Foucault, Michel. 1994. *The Order of Things: An Archaeology of Human Sciences*. New York: Vintage. Reprint edition.
- Fuke, Sawa, Masaki Ogino, and Minoru Asada. 2007. "Body Image Constructed from Motor and Tactile Images with Visual Information." *International Journal of Humanoid Robotics* 4 (2): 347–364.
- Gallagher, Shaun. 2000. "Philosophical Conceptions of the Self: Implications for Cognitive Science." *Trends in Cognitive Sciences* 4 (1): 14–21. <http://www.sciencedirect.com/science/article/pii/S1364661399014175>
- Hikita, Mai, Sawa Fuke, Masaki Ogino, Takashi Minato, and Minoru Asada. 2008. "Visual Attention by Saliency Leads Cross-Modal Body Representation." In *The 7th International Conference on Development and Learning*. New York: IEEE.
- Hinton, Geoffrey E., and Steven J. Nowlan. 1987. "How Learning Can Guide Evolution." *Complex Systems* 1:495–502.
- Hoffmann, Matej, Hugo Gravato Marques, Alejandro Hernandez Arieta, Hidenobu Sumioka, Max Lungarella, and Rolf Pfeifer. 2010. "Body Schema in Robotics: A Review." *IEEE Transactions on Autonomous Mental Development* 2 (4): 304–324.
- Hopson, Janet L. 1998. "Fetal Psychology." *Psychology Today* 31 (5): 44–64.
- Hubel, David H., and Torsten N. Wiesel. 1970. "The Period of Susceptibility to the Physiological Effects of Unilateral Eye Closure in Kittens." *Journal of Physiology* 206 (2): 419–436.
- Ishihara, Hisashi, Yuichiro Yoshikawa, and Minoru Asada. 2011. "Realistic Child Robot 'Affetto' for Understanding the Caregiver-Child Attachment Relationship that Guides the Child Development." In Vol. 2, *2011 IEEE International Conference on Development and Learning*, 1–5. New York: IEEE.
- Jacquey, Lisa, Gianluca Baldassarre, Vieri Giuliano Santucci, and J. Kevin O'Regan. 2019. "Sensorimotor Continuities as a Key Driver of Development: From Babies to Robots." *Frontiers in Neurorobotics* 13, article 98.
- Karmiloff-Smith, Annette. 1995. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.
- Kumar, Sanjeev, and Peter Bentley, eds. 2003. *On Growth, Form and Computers*. Cambridge, MA: Academic Press.
- Kuniyoshi, Yasuo, and Shinji Sangawa. 2006. "Early Motor Development from Partially Ordered Neural-Body Dynamics: Experiments with a Cortico-Spinal-Musculo-Skeletal Model." *Biological Cybernetics* 95 (6): 589–605.
- Lee, Mark H., Qinggang Meng, and Fei Chao. 2007. "Staged Competence Learning in Developmental Robotics." *Adaptive Behavior* 15 (3): 241–255.
- Lungarella, Max, Giorgio Metta, Rolf Pfeifer, and Giulio Sandini. 2003. "Developmental Robotics: A Survey." *Connection Science* 15 (4): 151–190.
- Mannella, Francesco, Vieri G. Santucci, Eszter Somogyi, Lisa Jacquey, Kevin J. O'Regan, and Gianluca Baldassarre. 2018. "Know Your Body through Intrinsic Goals." *Frontiers in Neurorobotics* 12:30.

- McKinney, Michael L., and Kenneth J. McNamara. 1991. *Heterochrony: The Evolution of Ontogeny*. New York: Plenum Press.
- Meltzoff, Andrew N., and M. Keith Moore. 1977. "Imitation of Facial and Manual Gestures by Human Neonates." *Science* 98 (4312): 75–78.
- Meola, Valentina Cristina, Daniele Caligiore, Valerio Sperati, Loredana Zollo, Anna Lisa Ciancio, Fabrizio Taffoni, Eugenio Guglielmelli, and Gianluca Baldassarre. 2015. "Interplay of Rhythmic and Discrete Manipulation Movements during Development: A Policy-Search Reinforcement-Learning Robot Model." *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 152–170.
- Mori, Hiroki, and Yasuo Kuniyoshi. 2010. "A Human Fetus Development Simulation: Self-Organization of Behaviors through Tactile Sensation." In *The 9th International Conference on Development and Learning*, 82–87. New York: IEEE.
- Morse, Anthony, Tony Belpaeme, Angelo Cangelosi, and Caroline Floccia. 2011. "Modeling U Shaped Performance Curves in Ongoing Development." In *Proceedings of the Annual Meeting of the Cognitive Science Society*. Seattle: Cognitive Science Society.
- Morse, Anthony, Tony Belpaeme, Angelo Cangelosi, and Linda B. Smith. 2010. "Thinking with Your Body: Modelling Spatial Biases in Categorization Using a Real Humanoid Robot." In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 1362–1367. Seattle: Cognitive Science Society.
- Nagai, Yukie, Minoru Asada, and Koh Hosoda. 2006. "Learning for Joint Attention Helped by Functional Development." *Advanced Robotics* 20 (10): 1165–1181.
- Nagai, Yukie, Koh Hosoda, and Minoru Asada. 2003a. "How Does an Infant Acquire the Ability of Joint Attention? A Constructive Approach." In *Proceedings of the 3rd International Workshop on Epigenetic Robotics*, 91–98. Lund, Sweden: Lund University Cognitive Studies.
- Nagai, Yukie, Koh Hosoda, Akio Morita, and Minoru Asada. 2003b. "A Constructive Model for the Development of Joint Attention." *Connection Science* 15 (4): 211–229.
- Neisser, Ulric, ed. 1994. "The Self Perceived." In *Emory Symposia in Cognition*, 3–22. Cambridge: Cambridge University Press.
- Nolfi, Stefano, and Dario Floreano, eds. 2000. *Evolutionary Robotics*. Cambridge, MA: MIT Press.
- Nolfi, Stefano, Domenico Parisi, and Jeffrey L. Elman. 1994. "Learning and Evolution in Neural Networks." *Adaptive Behavior* 3 (1): 5–28.
- Ogino, Masaki, Masaaki Kikuchi, and Minoru Asada. 2006. "Active Lexicon Acquisition Based on Curiosity." In *The 5th International Conference on Development and Learning*. New York: IEEE.
- Oudeyer, Pierre-Yves, Verena V. Hafner, and Andrew Whyte. 2005. "The Playground Experiment: Task-Independent Development of a Curious Robot." In *Proceedings of the AAAI Spring Symposium on Developmental Robotics*, 42–47. Menlo Park, CA: AAAI Press.
- Oudeyer, Pierre-Yves, Frédéric Kaplan, and Verena V. Hafner. 2007. "Intrinsic Motivation Systems for Autonomous Mental Development." *IEEE Transactions on Evolutionary Computation* 11 (2): 265–286.
- Pfeifer, Rolf, and Josh C. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Piaget, Jean. 1952. *The Origins of Intelligence in Children*. New York: International Universities Press.
- Ridley, Matt. 2003. *Nature via Nurture: Genes, Experience, and What Makes Us Human*. New York: HarperCollins.
- Rucinski, Marek, Angelo Cangelosi, and Tony Belpaeme. 2011. "An Embodied Developmental Robotic Model of Interactions between Numbers and Space." In *Proceedings of the Annual Meeting of the Cognitive Science Society*. Seattle: Cognitive Science Society.
- Santucci, Vieri Giuliano, Gianluca Baldassarre, and Marco Mirolli. 2016. "GRAIL: A Goal-Discovering Robotic Architecture for Intrinsically-Motivated Learning." *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 214–231.
- Schlesinger, Matthew, Dima Amso, and Scott P. Johnson. 2007. "Simulating Infants' Gaze Patterns during the Development of Perceptual Completion." In *Proceedings of the Seventh International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund, Sweden: Lund University Cognitive Studies.
- Schmidhuber, Jürgen. 2010. "Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990–2010)." *IEEE Transactions on Autonomous Mental Development* 2 (3): 230–247.
- Seepanomwan, Kristiana, Daniele Caligiore, Angelo Cangelosi, and Gianluca Baldassarre. 2015. "Generalisation, Decision Making, and Embodiment Effects in Mental Rotation: A Neurobotic Architecture Tested with a Humanoid Robot." *Neural Networks* 72:31–47.
- Stanley, Kenneth, and Risto Miikkulainen. "A Taxonomy for Artificial Embryogeny." *Artificial Life* 9 (20032): 93–130.

- Stewart, John, Olivier Gapenne, and Ezequiel A. Di Paolo. 2010. *Enaction toward a New Paradigm for Cognitive Science*. Cambridge, MA: MIT Press.
- Sumioka, Hidenobu, Yuichiro Yoshikawa, and Minoru Asada. 2010. "Reproducing Interaction Contingency toward Open-Ended Development of Social Actions: Case Study on Joint Attention." *IEEE Transactions on Autonomous Mental Development* 2 (1): 40–50.
- Tani, Jun. 2016. *Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. Cambridge: Oxford University Press.
- Thelen, Esther, Gregor Schöner, Christian Scheier, and Linda B. Smith. 2002. "The Dynamics of Embodiment: A Field Theory of Infant Perseverative Reaching." *Behavioral and Brain Sciences* 24:1–86.
- Thelen, Esther, and Linda B. Smith. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Varela, Francisco J., Eleanor Rosch, and Evan Thompson. 1992. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
- Verbeek, Peter-Paul. 2011. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press.
- Vernon, David. 2010. "Enaction as a Conceptual Framework for Developmental Cognitive Robotics." *Paladyn* 1:89–98.
- Vieira Neto, Hugo, and Ulrich Nehmzow. 2007. "Real-Time Automated Visual Inspection Using Mobile Robots." *Journal of Intelligent and Robotic Systems* 49 (3): 293–307.
- Warneken, Felix, Frances Chen, and Michael Tomasello. 2006. "Cooperative Activities in Young Children and Chimpanzees." *Child Development* 77 (3): 640–663.
- Weng, Juyang, James McClelland, Alex Pentland, Olaf Sporns, Ida Stockman, Mriganka Sur, and Esther Thelen. 2001. "Autonomous Mental Development by Robots and Animals." *Science* 291 (5504): 599–600.
- Wilson, Margaret. 2002. "Six Views of Embodied Cognition." *Psychonomic Bulletin and Review* 9 (4): 625–636.
- Yamada, Yasunori, Hoshinori Kanazawa, Sho Iwasaki, Yuki Tsukahara, Osuke Iwata, Shigehito Yamada, and Yasuo Kuniyoshi. 2016. "An Embodied Brain Model of the Human Foetus." *Scientific Reports* 6:27893.
- Yamada, Yasunori, Hiroki Mori, and Yasuo Kuniyoshi. 2010. "A Fetus and Infant Developmental Scenario: Self-Organization of Goal-Directed Behaviors Based on Sensory Constraints." In *Proceedings of the Tenth International Conference on Epigenetic Robotics*, 142–152. Lund, Sweden: Lund University Cognitive Studies.
- Yoshikawa, Yuichiro, Hiroyoshi Kawanishi, Minoru Asada, and Koh Hosoda. 2002. "Body Scheme Acquisition by Cross Modal Map Learning among Tactile, Visual, and Proprioceptive Spaces." In *Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 181–184. Lund, Sweden: Lund University Cognitive Studies.
- Ziemke, Tom. 2001. "Are Robots Embodied?" In *First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund, Sweden.
- Zlatev, Jordan, and Christian Balkenius. 2001. "Why Epigenetic Robotics." In *First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund, Sweden: Lund University Cognitive Studies.

4 Evolutionary Robotics

Stefano Nolfi

4.1 Introduction

Evolutionary robotics (Nolfi and Floreano 2000; Nolfi et al. 2016; Nolfi, 2021) is a method that allows the creation of robots capable of developing the ability to perform one or more functions as a result of an adaptation process analogous to natural evolution.

Robots are considered to be autonomous artificial organisms that adapt in close interaction with the environment without human intervention. The role of the experimenter is limited to the specification of the fitness function—that is, the criteria used to evaluate the performance level of the robots—and to the specification of the characteristics of the robots that are not subjected to the adaptive process. The remaining characteristics are encoded in a vector of parameters (genotype) and evolved through an evolutionary algorithm (Rechenberg 1973; Goldberg and Holland 1988). In the majority of cases, the evolving robots are provided with neural network controllers. The connection weights of the network, which determine the behavior of the robot, are encoded in the genotype and evolved. Eventually, the architecture of the neural network (Stanley and Miikkulainen 2002; Durr, Mattiussi, and Floreano 2006) and/or the morphology of the robot can be encoded in the genotype and evolved (Sims 1994; Lipson and Pollack 2000; Auerbach and Bongard 2012; Hiller and Lipson 2012).

The evolutionary process is realized by creating an initial population of genotypes generated randomly and then repeating the following steps for a certain number of generations: 1) create a population of robots with the characteristics specified in the corresponding genotypes, 2) allow the robots to interact with their environment for a finite amount of time and calculate a scalar value (fitness) that rates the performance of each robot with respect to a given problem, and 3) create a new population of genotypes composed of copies with random variations of the genotypes of the fittest robots.

An important aspect to consider is that the utilization of a fitness function that rewards the robot for performing a given function—for example, foraging—can drive the development of several behavioral and cognitive capacities that are instrumental to the achievement of that function, such as avoiding obstacles and dangers, orienting and navigating in the environment, discriminating relevant objects, integrating sensory information over time and later using it to appropriately regulate the robot's behavior, and so on. The analysis of the

way in which these capacities are realized and integrated in evolving robots can provide valuable information from the perspective of modeling the organization and the development of similar capacities in natural systems.

Evolutionary robotics has been applied to the study of a wide range of phenomena, including embodied cognition, sensorimotor coordination, integration of behavioral and cognitive skills, social and collective behaviors, internal models, and interaction between evolution and learning. In the following sections, I will describe a few representative examples of the work conducted in these areas.

4.2 Evolving Bodies and Brains: Morphological Computation

The behavioral and cognitive skills of robots or animals are dynamical properties that unfold in time and arise from a large number of interactions between the agent's nervous system, body, and environment (Chiel and Beer 1997; see also chapter 11). The dynamical process originating from the interactions depends on the characteristics of the agent's body and brain. This implies that varying the characteristics of the body and/or of the brain can shape the dynamical process.

An example of behavior that can be realized by shaping the characteristics of the body or of the brain is walking on a declining plane. Indeed, it can be produced either by brainless robots with passive joints and carefully designed body morphologies (McGeer 1990; Collins et al. 2005) or by highly controlled robots lacking the morphological features of the former robots (Chestnutt et al. 2005). The term “morphological computation” (Pfeifer et al. 2006; Paul 2006; see also chapter 1) has been introduced to indicate processes performed by the body that otherwise would have to be performed by the brain. Solutions exploiting morphological computation are often advantageous in terms of energy efficiency and robustness with respect to alternative solutions (Pfeifer and Bongard 2006).

The possibility of adapting both the body plan and the control policy of robots permits the selection of solutions that are simpler and more effective within the spectrum of those available—that is, among solutions relying primarily on morphological computation or on control. Moreover, it permits the generation of solutions in which the morphological and control features are coadapted. Evolutionary robotics constitutes an ideal approach for adapting both the policy and the morphology of robots since it is a model-free method that does not make any assumption about the structure of the adaptive system. Moreover, unlike alternative model-free training methods, it permits the adaptation of any type of parameter, including a combination of qualitatively different parameters. The number of body parts forming the body of the robot, the relative position of these parts, the physical properties of each body part, and the characteristics of the joints among body parts can be encoded in the genotype and evolved together with the characteristics of the neural network of the robot. This is typically realized by using genotypes that encode growing rules, which determine how the initial “embryo” grows and differentiates, rather than using genotypes that directly encode the property of a fully formed robot.

In a pioneering work in this area, Sims (1994) demonstrated how artificial evolution can be used to evolve the morphology and the control policy of simulated creatures capable of swimming, walking, and grabbing objects while competing with other creatures. Lipson and

Pollack (2000) later used a similar approach to evolve simulated walking robots that are then manufactured using a three-dimensional printer and spare electronic components.

Since that time, this approach has been used for various purposes. For example, Long (2012) evolved the stiffness of artificial tails of swimming robots to investigate how backbones evolved in early vertebrates. By evolving robots in environments of varying complexity, Auerbach and Bongard (2012) showed how the complexity of the evolved morphology correlates with the complexity of the environment. For example, robots evolved to walk on irregular terrain develop morphologies that include appendages missing in robots evolved over flat terrain. Hiller and Lipson (2012) demonstrated how evolving robots made of cells with different material properties arranged in evolved topologies can produce a variety of locomotion behaviors. These behaviors originate from simple periodic expansion/contraction actions produced by some of the cells and from the physical interactions among the cells composing the robot body and among the cells and the environment. These simulated robots composed of multiple cells can then be transformed into artificial living creatures by assembling ectoderm and cardiac stem cells in the same three-dimensional spatial configuration (Kriegman et al. 2020). Remarkably, these artificial living creatures are able to locomote and to explore their aqueous environment autonomously for days.

4.3 Sensorimotor Coordination

In agents that are embodied and situated, the role of perception cannot be separated by that of action and vice versa. What an agent perceives is determined by what it does, and what an agent does can be determined by what the agent needs to perceive.

The existence of a close link between perception and action draws on a number of distinct traditions in philosophy, in psychology, and in the cognitive sciences. It is at the core of the ecological theory of perception developed by Gibson (1979) and of several other fundamental contributions (Arbib 1989; Varela, Thomson, and Rosh 1991; Maturana and Varela 1987; Thelen and Smith 1994; Berthoz 2000; O'Regan and Noë 2001; Noë 2004; Clark 1998, 1999). The coupling of the sensory and the motor process can be indicated with the term “sensorimotor coordination” (Dewey 1981 [1986]).

Evolutionary robotics constitutes an ideal framework for studying the role of sensorimotor coordination in the development of behavioral and cognitive skills. The first reason is that the evolutionary process leaves the evolving robots free to determine the way in which they achieve their adaptive goals. Consequently, the robots are free to coordinate their perceptual and action processes in ways that are functional to the achievement of their objectives. The second reason is that the evolutionary process is driven by a fitness measure that rates the overall performance of the robot—that is, the sum of rewards obtained over an extended evaluation period. This permits variations that enhance the coordination between the sensory and action process to be identified and retained regardless of whether the time interval between actions and associated rewards is immediate or delayed.

Indeed, sensorimotor coordination plays a crucial role in practically all experiments carried out by evolving robots. The first demonstration was reported in an experiment in which a wheeled robot provided with infrared sensors and situated in an arena surrounded by walls was evolved for the ability to find and remain near a cylindrical object (Nolfi 1996, 2005). Interestingly, the evolved robots did not solve the problem by internally processing

the experienced sensory states in order to discriminate the stimuli corresponding to walls and cylinders, a strategy that was actually challenging since the stimuli experienced near cylinders and walls strongly overlap in sensory space. They instead solved the task by reacting to the stimuli to produce behavioral attractors—that is, oscillatory behavior generated by alternating move-forward/move-backward and turn-left/turn-right actions, near cylinders but not near walls. In other words, they exploited the fact that the execution of the same actions has different perceptual consequences near walls or cylinders that can lead to the production of the two required differentiated behaviors. This experiment can be replicated with the Evorobotpy software tool available from <https://github.com/snolfi/evorobotpy> (see the instruction for running the ErDiscrim experiment in Nolfi 2021, chapter 13).

In an extended version of this experiment, in which the robot was provided with proprio-sensors that encoded the speed of the robot's wheels, Scheier, Pfeifer, and Kunyoshi (1998) observed the evolution of a qualitatively different sensorimotor strategy that exploits actions to self-select easy-to-interpret stimuli. In this case the evolved robots displayed a wall-following behavior near walls and cylinders of moving straight along the wall and turning around the cylinder, respectively. They then used the perceived offset between the speed of the left and right wheel to keep producing the wall-following behavior near cylinders and to move away from walls. In other words, the robots acted to later experience favorable sensory states. They displayed an initial behavior that enabled them to later experience two well-differentiated states on their proprioceptors near walls and cylinders.

Qualitatively similar solutions have been observed in more complex robots evolved for the ability to solve more challenging problems. This is the case, for example, of an experiment in which a simulated iCub robot (Sandini, Metta, and Vernon 2004) was evolved for the ability to discriminate spherical and ellipsoid objects on the basis of rough tactile information (Tuci, Massera, and Nolfi 2010). The robot was provided with fourteen motor neurons that encoded the torque produced by seven sets of antagonistic muscles controlling the seven degrees of freedom (DOFs) of the arm and of the wrist, two motor neurons that encoded the desired extension/flexion of the thumb and of the four fingers, and two motor neurons that indicated the category of the object (i.e., spherical or ellipsoid). The sensors of the robot included eight neurons that encoded the current angular position of the DOFs of the arm and of the wrist, five neurons that encoded the extension/flexion of the five corresponding fingers, and ten neurons that encoded the ten touch sensors located on the fingertips and on the palm. Touch sensors binarily encoded whether the corresponding part of the robot body collided with another body. The robots were rewarded for discriminating the shape of the objects experienced during multiple evaluation episodes. They were not rewarded for the production of any specific behaviors and consequently were left free to select behaviors that enabled and/or facilitated the discrimination problem.

The analysis of the evolved robots demonstrates that they did indeed develop manipulation behaviors that enabled them to experience stimuli allowing them to reliably discriminate the two types of objects despite the similarity of the objects' shapes and the limited resolution of the touch sensors. The categorization process involves three phases. In the first part, the robot manipulates the object by wrapping it with its fingers and by moving the object until a suitable hand/object posture is reached. The information contained in the tactile stimuli experienced during this phase increases and finally reaches a high value when a hand/object achieves a suitable posture, which remains almost stable in the remain-

ing part of the episode. During the second phase, the robot starts to produce a categorization answer, keeps producing fine manipulation actions, and keeps integrating the sensory information experienced by eventually reversing its categorization decision. This continues during the third phase, in which the categorization decision is no longer reversible.

The solutions discovered by the evolved robots thus fit the dynamical view of cognition elaborated by Spivey (2007). The extension of the categorization process over time enables the robot to experience useful stimuli and to integrate the conflicting evidence experienced over time in order to maximize the accuracy of the categorization decision.

4.4 On the Relation between Reactive and Cognitive Capabilities

Evolutionary robotics can also be used to study the relation and the integration between behavioral and cognitive capabilities.

As discussed above, morphological computation and sensorimotor coordination can be used to perform processes that the brain would otherwise have to perform. The exploitation of the interaction between the agent and the environment thus permits reliance on solutions that are simpler, from an internal-processing perspective, than solutions that do not rely on these properties. This opens up questions about the relationship between reactive and cognitive capabilities. Do they tend to interact in a synergetic or conflictual manner? And “is cognition truly seamless—implying a gentle, incremental trajectory linking fully embodied responsiveness to abstract thought and off-line reason? Or is it a patchwork quilt, with jumps and discontinuities and with very different kinds of processing and representations serving different needs?” (Clark 1999, 350).

Interesting evidence supporting a synergetic relation and a smooth incremental integration of reactive and cognitive capabilities has been reported in evolutionary experiments addressing the evolution of a robot selected for the ability to navigate in a double T-maze environment (figure 4.1; Carvalho and Nolfi 2016). The robot, which is initially located in an area at the bottom of the central corridor with a randomly varying position and orientation, should




Figure 4.1
The object-discriminating robot.

travel toward a target destination located at one of the four ends of the maze. The correct destination is marked by two green objects located in the central corridor. The robot should thus solve a time-delay problem in which the information experienced while it travels down the central corridor should later influence the direction in which the robot turns when it reaches the first and the second junction.

The analysis of evolving robots indicates that they solve the problem with a strategy that does not require them to store the information extracted from the green object in internal states, recognize the arrival at the first and at the second junction, or turn left or right on the basis of the internal states and of the junction. As shown in figure 4.2, the trajectories produced during different evaluation episodes first converge in the bottom portion of the central corridor and then diverge while the robot perceives the position of the green objects. The initial convergence enables the robot to reduce the differences caused by the varying initial positions and orientations. The divergent process allows the robot to enter into one of four separate basins of attraction of robot/environmental dynamics that bring the robot to the right destination—the destination that matches the relative position of the two green beacons.

The strategy displayed by evolved robots thus exploits a form of cognitive off-loading—that is, the possibility of off-loading an agent's future intention into the external environment (Gilbert 2015a, 2015b). More specifically, the robot off-loads the information experienced in the central corridor by assuming different positions and orientations with respect to the corridor and by then maintaining such positions/orientations. The relative position of the robot in the corridor is then used to turn appropriately left or right at the first and then at the second junction. The trajectories displayed in figure 4.2 are produced by a robot that has no memory. However, similar strategies are produced by robots with memory—that is, by robots provided with recurrent connections in their internal neurons. The possibility of off-loading information




Figure 4.2

Trajectories of a typical evolved robot postevaluated for three hundred episodes. The trajectories (*shown in magenta, blue, yellow, and cyan*) indicate those produced by the robot during episodes in which it should have navigated toward the destination with the corresponding color. The target destination is marked by the relative position of the two green objects located to the left or right of the central corridor.

in the environment is thus preferred to alternative solutions relying on internal processing independently from the availability of memory.

Interestingly, evolving robots subjected to position perturbations, such as being randomly moved left or right as a result of “gusts of wind” occurring from time to time, solve the problem by developing composite strategies that rely on cognitive off-loading to determine the motor trajectory and on memory to reenter the appropriate basin of attraction after a position perturbation. This and additional control experiments reported in Carvalho and Nolfi (2016) demonstrate how, at least in this domain, reactive strategies do not prevent but rather promote the development of cognitive capabilities. Moreover, they illustrate how the development of cognitive capacities does not lead to the elimination of preexisting reactive capacities but rather to their extension.

4.5 Social and Collective Behavior

In the previous section, we limited our analysis to individual behaviors—to the evolution of robots placed in an environment that does not include other robots. The evolutionary method, however, can also be applied to evolve social behaviors. This can be done simply by situating the evolving robots in environments containing other robots.

This scenario has been used to study the conditions that support the evolution of cooperative behavior. As expected, cooperative behavior readily emerges when a group of interacting robots is formed by genetically related individuals (e.g., individuals possessing identical genotypes) or when selection operates at the level of the colony or swarm (Floreano et al. 2007). When instead the individuals forming the colony are not genetically related and selection operates at the level of individuals, the evolutionary process leads to a dynamic in which cooperation periodically emerges and extinguishes (Mitri, Floreano, and Keller 2009).

The evolution of genetically related robots readily produces self-organizing properties—that is, the spontaneous formation of spatial, temporal, or spatiotemporal structures or functions that emerge from local interactions among individual robots and that are robust with respect to environmental variations (Camazine et al. 2001; see also chapter 5). For example, Sperati, Trianni, and Nolfi (2011) conducted experiments in which a population of wheeled robots was evolved for the ability to forage. The evolving robots developed an ability to arrange themselves in dynamic chains that enabled the colony to efficiently navigate between a nest and a foraging area. These dynamic chains, which self-sustain in the presence of perturbations, allow robots with limited individual sensory capacities to efficiently navigate to the right destination by discovering and storing information on the location of the relevant environmental areas at the level of the colony. Another example of self-organized behavior has been observed in a population of robots capable of self-assembly—in this case, by physically attaching together—to master problems that cannot be solved by individual robots. Robots evolved for the ability to move while attached developed an ability to negotiate a common direction of motion and to keep moving along that direction by compensating for misalignments originating during motion (Baldassarre et al. 2007). Also in this case, the ability to coordinate and to cooperate was robust with respect to variations in the environmental conditions. Indeed, evolved robots

were capable of coordinating independently from the configuration in which they were assembled. Moreover, robots evolved in specific environmental conditions demonstrated the ability to generalize their skills to new environmental conditions. Such generalization capacity included the ability to display new behaviors adapted to the new experienced conditions. For example, robot swarms evolved in an environment with no obstacles demonstrated an ability to avoid obstacles and to rearrange their shape to pass through narrow passages when situated in a mazelike environment with obstacles (Nolfi 2009).

The evolution of collective behavior in robots can also lead to the emergence of task specialization—that is, to individuals capable of assuming different complementary roles that increase the efficacy of the group (Ferrante et al. 2015; Pagliuca and Nolfi 2018).

The evolution of robots selected for the ability to solve a problem that benefits from cooperation has also been used to study the evolution of communication and language (Cangelosi and Parisi 2002; Nolfi and Mirolli 2010; see also chapter 20). In a series of experiments reported in De Greef and Nolfi (2010), the authors analyzed the origin and complexification of the communication system displayed by evolving robots across generations and the origin and transformation of the meaning associated with communication signals. These analyses indicate that the development of communication capabilities is strongly interlinked with the evolution of other capabilities. Robots need to develop appropriate behaviors to access and/or generate the information to be communicated and to react appropriately to detected signals. Interestingly, the development of communication skills scaffolds the development of behavioral skills and vice versa. This leads to the development of integrated capabilities and to a progressive complexification of robots' skills (Nolfi 2013).

Finally, evolutionary robotics experiments have been used to explain why reciprocity, the reciprocal exchange of episodes of help between two partners, is rare in nature (André and Nolfi 2016). This fact contrasts with the predictions generated by game theoretic models that reciprocity should evolve easily (Axelrod and Hamilton 1981). As shown by André and Nolfi (2016), these game theoretic models' predictions are in error because these methods do not model the mechanisms underlying the generation of behavior, a limitation that does not affect evolutionary robotics models. Indeed, the experiments carried out by evolving robots predict correctly that reciprocity is unlikely to evolve, due to the numerous neutral mutations required to generate a reciprocator behavior from individuals that do not reciprocate.

Another line of research has investigated the evolution of social behaviors in competing scenarios—for example, the evolution of a population of robots with conflicting interests. The coevolution of competing species such as predator and prey might favor the synthesis of evolutionary innovations. Indeed, “an adaptation in one lineage (e.g., predators) may change the selection pressure on another lineage (e.g., prey), giving rise to a counteradaptation. If this occurs reciprocally, an unstable runaway escalation of ‘arm races’ may result” (Dawkins and Krebs 1979, 489; Rosin and Belew 1997). In other words, adaptations on one side call for counteradaptations on the other side, and the counteradaptations call for more counteradaptations, and so on, thus producing an escalation process. Moreover, the concurrent evolution of the agents and of the learning environment can lead to a spontaneous, progressive complexification of the adaptive problem. That is to say, a pedagogically sound training process can be produced in which progress in one population

is accompanied by a gradual complexification of the adaptive task caused by parallel progress in the competing population (Rosin and Belew 1997).

Evolutionary experiments performed by evolving predator and prey robots (Cliff and Miller 1995; Nolfi and Floreano 1998) showed that co-evolution does indeed lead to “arms races” that produce a progressive complexification during the initial generations. The evolutionary dynamics, however, later converge in a limit-cycle dynamic in which progress against current competitors (local progress) is accompanied by retrogression with respect to ancient or future competitors. Cycling dynamics of this type were found in natural evolution in a population of side-blotched lizards (*Uta stansburiana*) by Sinervo and Lively (1996) and in *Daphnia* and associated parasites conserved in lake sediment (Decaestecker et al. 2007). More recently, Simione and Nolfi (2017, 2019) showed how long-term global progress can be produced in controlled ecological conditions—that is, in experiments in which the evolving populations are divided into subgroups that normally interact with specific subgroups of the competing population and only occasionally with the remaining competitors.

4.6 Evolution, Development, and Learning

The basic evolutionary method illustrated in the introduction can be extended to incorporate development and learning. In the basic method, the process that maps a genotype into a robot is completed before the robot starts to interact with its environment. In other words, robots are born as fully formed individuals. In extended evolutionary methods, by contrast, the developmental process continues during the period in which the robot interacts with its environment.

A model described in Bongard (2011), in which the evolving robots developed from an anguilliform morphology to a legged morphology while they interacted with the external environment, provides an example. The comparison with control experiments, in which the robots did not transition through the anguilliform body plan, indicates that morphological change accelerates the evolution of robust walking behaviors. A second example is given by a series of experiments reported in Kriegman, Cheney, and Bongard (2018) in which soft robots with developmental morphology were evolved for the ability to move over a surface. The analysis of the interaction between the evolutionary and developmental processes in these experiments enabled the authors to highlight an unknown aspect of genetic assimilation—namely, that the traits that render the agents robust to changes in other traits have a greater probability of becoming genetically assimilated in successive generations than traits that are less robust to genetic variations.

A model in which the brains of the robots keep developing while the robots operate in their environment was studied in Nolfi, Miglino, and Parisi (1994). In this model, the evolving robots were provided with neuron axons that grew and branched by establishing connections with other neurons while the robots operated in the environment. As with real nervous systems, the growth process of axons is influenced both by the activity patterns of the single neurons and by genetic factors (Purves 1994; Quartz and Sejnowski 1997). This leads to the evolution of robots capable of developing brains adapted to the environment in which they are situated—for example, to robots that might or might not develop

a brain area dedicated to processing light and in which development of the area is triggered by the exposure to light (Nolfi, Miglino, and Parisi 1994).

Other works have investigated the combination of evolution and learning (Nolfi and Floreano 1999). In these models the topology of the neural network was fixed, but the connection weights varied while the robots interacted with the environment on the basis of an unsupervised (Floreano, Durr, and Mattiussi 2008), self-supervised (Nolfi and Parisi 1993), or reinforcement-learning algorithm (Schembri, Mirolli, and Baldassarre 2007). The combination of evolution and learning enables evolving robots to adapt to environmental variations that occur within generations. For example, it enables predator robots to modify their behavior on the fly while interacting with a prey robot to display the strategy that is effective against the current encountered prey (Floreano and Nolfi 1997).

4.7 Internal Models

Evolutionary robotics is a model-free approach, a method that permits the robots to develop behavioral and cognitive skills from scratch without the need to rely on a model of the external environment and/or the robot's own self. However, the abilities that the robots develop during their adaptation can include the ability to build and use a model of their own body, a model of the external environment, and/or a forward model that allows the consequences of the robots' actions to be predicted.

Bongard, Zykov, and Lipson (2006) give an example of a robot capable of acquiring a model of its own body. In this work, a physical robot was equipped with an onboard simulator that it used to continually evolve a model of itself. The model consisted of a three-dimensional description of the robot's own body that enabled it to predict the perceptual effects of the actions it could execute without actually performing them. The robot then used the model to cope with damages, such as the mechanical separation of a leg. This was realized by 1) using the offset between the actual and predicted consequences of actions to diagnose the damage, 2) updating the model of the robot's own body to reduce the offset between the predicted and actual consequences of the robot's action, 3) evolving a new control policy capable of operating effectively with the damaged body by using a mental simulation, and 4) using the new control policy to keep operating effectively despite the damage. The availability of the world model thus permits the evolution of a compensatory policy by using the imagined effect of variations of the current policy (mental simulation) as a proxy for the actual effect of variations.

Cully et al. (2015) showed how the ability to recover from damages or faults can be speeded up by learning a behavior-performance map that encodes the correlation between the value of the connection weights and the value of fitness. The map can then be used to introduce mutations that have a higher chance of producing improvements with respect to random mutations.

Gigliotta, Pezzulo, and Nolfi (2011) demonstrated how a robot subjected to sensory deprivation can evolve the ability to react appropriately to sensory stimuli and to self-generate states functionally equivalent to sensory stimuli during sensory deprivation phases in which stimuli are not available. The behavior consists of moving the robot's eye to foveate consecutive portions of the image located over a circular trajectory. In normal phases, the robot can

determine the movement of the eye on the basis of the current perceived color. During blind phases, the robot should use self-generated internal states as proxies for missing sensory states. The analysis of the evolved robots indicates that the problem is not solved by generating states that match the missing sensory states. Rather, it is realized by generating internal states that elicit the appropriate movements but are not necessarily similar to the states that would be experienced in normal conditions.

Finally, Ha and Schmidhuber (2018) demonstrated how agents that determine their actions on the basis of features extracted from the sensory states, by a neural network trained with a self-supervised learning algorithm, outperform agents that determine their actions directly on the basis of the features encoded in sensory states. The problem considered consists of learning to drive in a car-racing environment called CarRacing-v0 (Brockman et al. 2016). The learning agent receives an image containing a top-down view of the car and the environment as input. The features are extracted by 1) a variational autoencoder network (Kingma and Welling 2013; Rezende et al. 2014) trained with the ability to encode perceived images in compact representations that can be used to reconstruct the original image and 2) a long short-term memory (LSTM) network (Hochreiter and Schmidhuber 1997) trained to predict the compressed state of the next perceived image on the basis of the compressed state of the current image and of the action the agent is going to perform. These two networks are pretrained using the images collected by the agent during several evaluation episodes in which the agent moves by performing random actions. The neural network controller of the evolving agents, which receives as input the internal state extracted by the sensors from the two pretrained networks described above, is evolved by using a standard evolutionary method for the ability to drive the car. In a second experiment performed by using the *VizDoom* game problem (Kempka et al. 2016), the authors showed that the autoencoder and LSTM prediction network described above can be used to evolve the agents in virtual worlds imagined by the agents themselves. The solutions evolved in these imagined worlds can then be successfully used to control the agent of a real *VizDoom* game.

4.8 Evolution as a Form of Learning

The evolutionary method can also be used to model ontogenetic learning (Schlesinger 2004). This is because the evolutionary algorithm constitutes one of the simplest yet most effective ways to evolve an embodied neural network through a trial-and-error process based on distal rewards. An example is illustrated in experiments in which an iCub humanoid robot (Sandini, Metta, and Vernon 2004) trained through an evolutionary method develops reaching and grasping skills analogous to those displayed by human infants from two to eighteen months of age (Savastano and Nolfi 2013). During this period, infants display a first transition from sweeping and unsuccessful arm movements to primitive, imprecise reaching and grasping behaviors and then a second transition leading to integrated and effective reaching and grasping behaviors (Konczak et al. 1995; Konczak, Borutta, and Dichgans 1997; Konczak and Dichgans 1997; von Hofsten and Rönnqvist 1993; Spencer and Thelen 2000).

As illustrated in figure 4.3 (*left*), the robot is set in an upright position in front of a suspended object. This setting is similar to that used by Hofsten (1982) to analyze the development of



Figure 4.3

The simulated setting (*left*) is derived from experiments carried out on infants (*center and right*) by von Hofsten (1982).

reaching and grasping behavior in infants (figure 4.3, *center and right*). The training of the robot is realized in three phases: 1) a prereaching phase in which the robot has simple prewired reflex behaviors, low visual acuity, and an immature nervous system; 2) a gross-reaching phase in which the robot has improved visual acuity and matured cortical areas; and 3) a fine-reaching phase in which the robot has access to perceptual information that encodes the relative position of the object with respect to the hand.

The analysis of the experiments shows that the lack of internal neural resources during the prereaching phase has an adaptive role (i.e., channels the developmental process toward better solutions during the gross-reaching phase) and a bias role (i.e., represents a necessary condition for the emergence of the exploratory motor-babbling behavior). This suggests that the later involvement of cortical areas (Martin 2005) can play an adaptive role in humans and might have evolved to accomplish this function. Moreover, analysis of the behavior displayed by the robots during the course of the training process shows that the following phenomena observed in infants originate spontaneously: 1) a reduced use (freeze) of the distal DOFs of the arm of the robot during the prereaching phase, 2) an exploratory (motor-babbling) behavior during the prereaching phase, and 3) a temporal regression of the reaching capabilities at the onset of the fine-reaching phase. The fact that these qualitative variations emerge spontaneously indicates that they do not necessarily reflect the presence of additional specific maturational constraints. They can be the manifestation of a general self-structuring process that operates by temporarily reducing the complexity of the motor space, of the sensory space, and of the relevant task space, respectively.

In contrast to reinforcement-learning algorithms (Sutton and Barto 2018) that represent the most common choice to model trial-and-error learning, evolutionary algorithms present advantages and drawbacks. The advantages include the possibility of adapting all the characteristics of the robot, including the robot's morphology and the architecture of the robot's neural network and the ability to operate well in the presence of sparse reward. Reinforcement-learning algorithms, on the other hand, are generally more sample efficient.

The development of new evolutionary algorithms that operate by estimating the local gradient (Hansen and Ostermeier 2001) and eventually rely on stochastic gradient optimizers to vary the adaptive parameters (Salimans et al. 2017) makes the usage of evolutionary methods even more attractive. Indeed, although these gradient-ascent methods can also operate on populations that include multiple parents, they are typically used with popula-

tions composed of a single parent producing several offspring. The evaluation of the offspring is used to estimate the local gradient, which in turn is used to vary the parameters of the parent. This implies that, as in ontogenetic learning, the adaptation process is realized by varying the parameters of a single individual.

As demonstrated by Salimans et al. (2017), modern evolutionary methods represent a scalable alternative to the state-of-the-art reinforcement-learning algorithm (Schulman et al. 2015, 2017). Indeed, they can be used to adapt neural network controllers with millions of parameters by achieving results that are competitive with reinforcement-learning methods. The results have been collected on state-of-the-art benchmarking problems: the Mujoco control problems that require controlling articulated robots (Todorov, Erez, and Tassa 2012) and the Atari games that require controlling game players that receive as input the images of the console (Bellemare et al. 2013).

4.9 Conclusion

Evolutionary robotics is not only a method for automatic robot development inspired by biology but also a tool for investigating open questions concerning natural systems such as, for example, the role of embodiment in cognition, the origins of symbolic communication, the relation between behavioral and cognitive capacities, and the mechanisms supporting the development of cooperative behaviors.

Despite initial skepticism demonstrated by representatives of mainstream disciplines and even by pioneers of the approach (Matarić and Cliff 1996), over the years an increasing number of researchers from a wide range of disciplines have adopted the method. The richness and fecundity of the approach combined with the novel opportunities granted by recent methodological progress suggest that it will continue to play an important role in the future.

Readers interested in acquiring hands-on knowledge on evolutionary robotics can access freely available tools that permit the replication of standard experiments and the design of new experiments (see Auerbach et al. 2014; Massera et al. 2014; Nolfi 2021; see also <https://github.com/snolfi/evorobotpy>).

Additional Reading and Resources

- A recent review of the field: Nolfi, S., J. Bongard, P. Husbands, and D. Floreano. 2016. “Evolutionary Robotics.” In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1423–1541. 2nd ed. Berlin: Springer Verlag.
- An article that illustrates in more detail the complex adaptive system nature of behavior and cognition in embodied agents: Nolfi, S. 2009. “Behavior and Cognition as a Complex Adaptive System: Insights from Robotic Experiments.” In *Handbook of the Philosophy of Science. Volume 10: Philosophy of Complex Systems*, edited by C. Hooker. General editors: Dov M. Gabbay, Paul Thagard, and John Woods. San Diego: Elsevier.
- A more detailed review of the field: Nolfi, S., and D. Florean, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press, 2000.

- Evorobotpy (Nolfi 2021; <https://github.com/snolfi/evorobotpy2>) is a simple and well-documented tool that can be used to perform evolutionary robotics experiments. The associated documentation (Nolfi 2021, chap. 13) includes tutorials and exercises.
- Farsa (Massera et al. 2014; <https://sourceforge.net/projects/farsa/>) is another software tool that can be used to conduct evolutionary robotics experiments.

References

- André, Jean-Baptiste, and Stefano Nolfi. 2016. “Evolutionary Robotics Simulations Help Explain Why Reciprocity Is Rare in Nature.” *Scientific Reports* 6:32785.
- Arbib, Michael A. 1989. *The Metaphorical Brain: Vol. 2, Neural Networks and Beyond*. New York: Wiley.
- Auerbach, Joshua, Deniz Aydin, Andrea Maesani, Przemyslaw Kornatowski, Titus Cieslewski, Grégoire Heitz, Pradeep Fernando, Ilya Loshchilov, Ludovic Daler, and Dario Floreano. 2014. “RoboGen: Robot Generation through Artificial Evolution.” In *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, edited by H. Sayama, J. Reiffel, S. Risi, R. Doursat, and H. Lipson. New York: MIT Press.
- Auerbach, Joshua E., and Joshua C. Bongard. 2012. “On the Relationship between Environmental and Morphological Complexity in Evolved Robots.” In *Proceedings of the 14th International Conference on Genetic and Evolutionary Computation*, 521–528. ACM Press.
- Axelrod, Robert, and William Donald Hamilton. 1981. “The Evolution of Cooperation.” *Science* 211:1390–1396.
- Baldassarre, Gianluca, Vito Trianni, Michael Bonani, Francesco Mondada, Marco Dorigo, and Stefano Nolfi. 2007. “Self-Organised Coordinated Motion in Groups of Physically Connected Robots.” *IEEE Transactions on Systems, Man, and Cybernetics* 37 (1): 224–239.
- Bellemare, Marc G., Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. “The Arcade Learning Environment: An Evaluation Platform for General Agents.” *Journal of Artificial Intelligence Research* 47:253–279.
- Bernstein, Nikolai. 1967. *The Co-ordination and Regulation of Movements*. Oxford: Pergamon Press.
- Berthoz, Alain. 2000. *The Brain’s Sense of Movement*. Cambridge, MA: Harvard University Press.
- Bongard, Joshua C. 2011. “Morphological Change in Machines Accelerates the Evolution of Robust Behavior.” *Proceedings of the National Academy of Science* 108:1234–1239.
- Bongard, Josh, Victor Zykov, and Hod Lipson. 2006. “Resilient Machines through Continuous Self-Modelling.” *Science* 5802 (314): 1118–1121.
- Brockman, Greg, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. “OpenAI Gym.” ArXiv: 1606.01540.
- Brunyé, Tad T., Aaron L. Gardony, Amanda Holmes, and Holly A. Taylor. 2018. “Spatial Decision Dynamics during Wayfinding: Intersection Prompt the Decision-Making Process.” *Cognitive Research: Principles and Implications* 8:13.
- Camazine, Scott, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Eric Bonabeau, and Guy Theraula. 2001. *Self-Organization in Biological Systems*. Princeton, NJ: Princeton University Press.
- Cangelosi, Angelo, and Domenico Parisi. 2002. *Simulating the Evolution of Language*. Berlin: Springer Verlag.
- Carvalho, Jônata Tyska, and Stefano Nolfi. 2016. “Cognitive Offloading Does Not Prevent but Rather Promotes Cognitive Development.” *PLoS One* 11 (8): e0160679.
- Chestnutt, Joel, Manfred Lau, German Cheung, James Kuffner, Jessica Hodgins, and Takeo Kanade. 2005. “Footstep Planning for the Honda ASIMO Humanoid.” In *Proceedings of the IEEE International Conference on Robotics and Automation*. New York: IEEE.
- Chiel, Hillel J., and Randall D. Beer. 1997. “The Brain Has a Body: Adaptive Behavior Emerges from Interactions of Nervous System, Body and Environment.” *Trends in Neuroscience* 20:553–557.
- Clark, Andy. 1998. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, Andy. 1999. “An Embodied Cognitive Science?” *Trends in Cognitive Sciences* 3:345–351.
- Cliff, Dave, and Geoffrey F. Miller. 1995. “Tracking the Red Queen: Measurement of Adaptive Progress in Coevolutionary Simulations.” In *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, edited by F. Moran, A. Moreno, J. J. Merelo, and P. Chacon. Berlin: Springer Verlag.
- Collins, Steve, Andy Ruina, Russ Tedrake, and Martijn Wisse. 2005. “Efficient Bipedal Robots Based on Passive-Dynamic Walkers.” *Science* 307 (5712): 1082–1085.

- Cully, Antoine, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. "Robots That Can Adapt Like Animals." *Nature* 521:503–507.
- Dawkins, Richard, and John Richard Krebs. 1979. "Arms Races between and within Species." *Proceedings of the Royal Society of London B* 205:489–511.
- Decaestecker, Ellen, Sabrina Gaba, Joost A. M. Raeymaekers, Robby Stoks, Liesbeth Van Kerckhoven, Dieter Ebert, and Luc De Meester. 2007. "Host-Parasite 'Red Queen' Dynamics Archived in Pond Sediment." *Nature* 450 (6): 870–874.
- De Greef, Joachim, and Stefano Nolfi. 2010. "Evolution of Implicit and Explicit Communication in a Group of Mobile Robots." In *Evolution of Communication and Language in Embodied Agents*, edited by S. Nolfi and M. Mirolli. Berlin: Springer Verlag.
- Dewey, John. 1981 [1986]. "The Reflex Arc in Psychology." *Psychological Review* 3:357–370. Reprinted in *The Philosophy of John Dewey*, edited by J. J. McDermott, 136–148. Chicago: University of Chicago Press, 1986.
- Dürr, Peter, Claudio Mattiussi, and Dario Floreano. 2006. "Neuroevolution with Analog Genetic Encoding." In *Proceedings of the Ninth Conference on Parallel Problem Solving from Nature*. Vol. 9. Berlin: Springer.
- Ferrante, Eliseo, Ali Emre Turgut, Edgar Duéñez-Guzmán, Marco Dorigo, and Tom Wenseleers. 2015. "Evolution of Self-Organized Task Specialization in Robot Swarms." *PLoS Computational Biology* 11 (8): e1004273.
- Floreano, Dario, Peter Dürr, and Claudio Mattiussi. 2008. "Neuroevolution: From Architectures to Learning." *Evolutionary intelligence* 1 (1): 47–62.
- Floreano, Dario, Sara Mitri, Stéphane Magnenat, and Laurent Keller. 2007. "Evolutionary Conditions for the Emergence of Communication in Robots." *Current Biology* 17:514–519.
- Floreano, Dario, and Stefano Nolfi. 1997. "Adaptive Behavior in Competing Co-evolving Species." In *Proceedings of the Fourth Conference on Artificial Life*, edited by P. Husband and I. Harvey, 378–387. Cambridge, MA: MIT Press.
- Gigliotta, Onofrio, Giovanni Pezzulo, and Stefano Nolfi. 2011. "Evolution of a Predictive Internal Model in an Embodied and Situated Agent." *Theory in Biosciences* 130 (4): 259–276.
- Gilbert, Sam J. 2015a. "Strategic Offloading of Delayed Intentions into the External Environment." *Quarterly Journal of Experimental Psychology* 68 (5): 971–992.
- Gilbert, Sam J. 2015b. "Strategic Use of Reminders: Influence of Both Domain-General and Task-Specific Metacognitive Confidence, Independent of Objective Memory Ability." *Consciousness and Cognition* 33:245–260.
- Gipson, James J. 1979. *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Lawrence Erlbaum.
- Goldberg, David E., and John Henry Holland. 1988. "Genetic Algorithms and Machine Learning." *Machine Learning* 3 (2–3): 95–99.
- Ha, David, and Jürgen Schmidhuber. 2018. "World Models." arXiv: 1803.10122.
- Hansen, Nikolaus, and Andreas Ostermeier. 2001. "Completely Derandomized Self-Adaptation in Evolution Strategies." *Evolutionary Computation* 9 (2): 159–195.
- Hiller, Jonathan, and Hod Lipson. 2012. "Automatic Design and Manufacture of Soft Robots." *IEEE Transactions on Robotics* 28 (2): 457–466.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–1780.
- Iida, F., and R. Pfeifer. 2005. "Morphological Computation: Connecting Body, Brain and Environment." *Japanese Scientific Monthly* 58 (2): 48–54.
- Kempka, Michał, Marek Wydmuch, Grzegorz Runc, Jakub Toczek, and Wojciech Jaśkowski. 2016. "Vizdoom: A Doom-Based AI Research Platform for Visual Reinforcement Learning." In *IEEE Conference on Computational Intelligence and Games*, 341–348. New York: IEEE.
- Kingma, Diederik P., and Max Welling. 2013. "Auto-Encoding Variational Bayes." arXiv: 1312.6114v10.
- Konczak, Jürgen, Maïke Borutta, and Johannes Dichgans. 1997. "The Development of Goal-Directed Reaching in Infants. II. Learning to Produce Task-Adequate Patterns of Joint Torque." *Experimental Brain Research* 113 (3): 465–474.
- Konczak, Jürgen, Maïke Borutta, Helge Topka, and Johannes Dichgans. 1995. "The Development of Goal-Directed Reaching in Infants: Hand Trajectory Formation and Joint Torque Control." *Experimental Brain Research* 106 (1): 156–168.
- Konczak, Jürgen, and Johannes Dichgans. 1997. "The Development toward Stereotypic Arm Kinematics during Reaching in the First 3 Years of Life." *Experimental Brain Research* 117 (2): 346–354.
- Kriegman, Sam, Douglas Blackiston, Michael Levin, and Josh Bongard. 2020. "A Scalable Pipeline for Designing Reconfigurable Organisms." *Proceedings of the National Academy of Sciences* 117 (4): 1853–1859.

- Kriegman, Sam, Nick Cheney, and Josh Bongard. 2018. "How Morphological Development Can Guide Evolution." *Scientific Reports* 8:13934.
- Lipson, Hod, and Jordan B. Pollack. 2000. "Automatic Design and Manufacture of Artificial Lifeforms." *Nature* 406: 974–978.
- Long, John. 2012. *Darwin's Devices: What Evolving Robots Can Teach Us about the History of Life and the Future of Technology*. New York: Basic Books.
- Massera, Gianluca, Tomassino Ferrauto, Onofrio Gigliotta, and Stefano Nolfi. 2014. "Designing Adaptive Humanoid Robots through the FARSA Open-Source Framework." *Adaptive Behavior* 22 (3): 255–265.
- Martin, John H. 2005. "The Corticospinal System: From Development to Motor Control." *Neuroscientist* 11 (2): 161–173.
- Matarić, Maja, and Dave Cliff. 1996. "Challenges in Evolving Controllers for Physical Robots." *Robotics and Autonomous Systems* 19 (1): 67–83.
- Maturana, Humberto R., and Francisco J. Varela. 1987. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambhala.
- McGeer, Tad. 1990. "Passive Dynamic Walking." *International Journal of Robotics Research* 9 (2): 62–82.
- Mitri, Sara, Dario Floreano, and Laurent Keller. 2009. "The Evolution of Information Suppression in Communicating Robots with Conflicting Interests." *PNAS* 106 (37): 15786–15790.
- Noë, Alva. 2004. *Action in Perception*. Cambridge, MA: MIT Press.
- Nolfi, Stefano. 1996. "Adaptation as a More Powerful Tool than Decomposition and Integration." In *Proceedings of the Workshop on Evolutionary Computing and Machine Learning, 13th International Conference on Machine Learning*, edited by T. Fogarty and G. Venturini. Bari, Italy: University of Bari Aldo Moro.
- Nolfi, Stefano. 2005. "Categories Formation in Self-Organizing Embodied Agents." In *Handbook of Categorization in Cognitive Science*, edited by H. Cohen and C. Lefebvre. Oxford: Elsevier.
- Nolfi, Stefano. 2009. "Behavior and Cognition as a Complex Adaptive System: Insights from Robotic Experiments." In Vol. 10, *Handbook of the Philosophy of Science: Philosophy of Complex Systems*, edited by C. Hooker. San Diego: Elsevier.
- Nolfi, Stefano. 2013. "Emergence of Communication and Language in Evolving Robots." In *New Perspectives on the Origins of Language*, edited by C. Lefebvre, B. Comrie, and H. Cohen, 533–554. Amsterdam: John Benjamins.
- Nolfi, Stefano. 2021. *Behavioral and Cognitive Robotics: An Adaptive Perspective*. Roma, Italy: Institute of Cognitive Sciences and Technologies, National Research Council (CNR-ISTC). <https://bacrobotics.com/>.
- Nolfi, Stefano, Josh Bongard, Phil Husbands, and Dario Floreano. 2016. "Evolutionary Robotics." In *Handbook of Robotics*, 2nd ed., edited by B. Siciliano and O. Khatib. Berlin: Springer Verlag.
- Nolfi, Stefano, and Dario Floreano. 1998. "Co-evolving Predator and Prey Robots: Do 'Arm Races' Arise in Artificial Life?" *Artificial Life* 4 (4): 311–335.
- Nolfi, Stefano, and Dario Floreano. 1999. "Learning and Evolution." *Autonomous Robots* 7 (1): 89–113.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
- Nolfi, Stefano, Orazio Miglino, and Domenico Parisi. 1994. "Phenotypic Plasticity in Evolving Neural Networks." In *Proceedings of the International Conference from Perception to Action*, edited by D. P. Gaussier and J-D. Nicoud, 146–157. Los Alamitos, CA: IEEE Computer Society Press.
- Nolfi, Stefano, and Marco Mirolli. 2010. *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.
- Nolfi, Stefano, and Domenico Parisi. 1993. "Auto-teaching: Networks That Develop Their Own Teaching Input." In *Proceedings of the Second European Conference on Artificial Life*, edited by J. L. Deneubourg, H. Bersini, S. Goss, G. Nicolis, and R. Dagonnier, 845–862. Brussels: Université Libre de Bruxelles.
- O'Regan, J. K. and A. Noë. 2001. "A Sensorimotor Account of Vision and Visual Consciousness." *Behavioral and Brain Sciences* 24 (5): 939–973.
- Pagliuca, Paolo, and Stefano Nolfi. 2018. "Robust Optimization through Neuroevolution." *PLoS One* 14 (3): e0213193.
- Paul, Chandana. 2006. "Morphological Computation: A Basis for the Analysis of Morphology and Control Requirements." *Robotic Autonomous Systems* 54 (8): 619–630.
- Pfeifer, Rolf, and Joshua C. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, Rolf, Fumiya Iida, and Gabriel Gómez. 2006. "Morphological Computation for Adaptive Behavior and Cognition." In *International Congress Series*, Vol. 1291, 22–29. San Diego: Elsevier Press.

- Purves, Dale. 1994. *Neural Activity in the Growth of the Brain*. Cambridge: Cambridge University Press.
- Quartz, Steven R., and Terrence J. Sejnowski. 1997. "The Neural Basis of Cognitive Development: A Constructivist Manifesto." *Behavioral and Brain Science* 4:537–555.
- Rechenberg, Ingo. 1973. *Evolutionsstrategie—Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Stuttgart: Frommann-Holzboog.
- Rezende, Danilo Jimenez, Shakir Mohamed, and Daan Wierstra. 2014. "Stochastic Backpropagation and Approximate Inference in Deep Generative Models." In Vol. 32, *Proceedings of the 31st International Conference on Machine Learning*. Red Hook, NY: JMLR.
- Rosin, Christopher D., and Richard K. Belew. 1997. "New Methods for Competitive Coevolution." *Evolutionary Computation* 5 (1): 1–29.
- Salimans, Tim, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. "Evolution Strategies as a Scalable Alternative to Reinforcement Learning." ArXiv:1703.03864v2.
- Sandini, Giulio, Giorgio Metta, and David Vernon. 2004. "Robotcub: An Open Framework for Research in Embodied Cognition." *International Journal of Humanoid Robotics* 8 (2): 18–31.
- Savastano, Piero, and Stefano Nolfi. 2013. "A Robotic Model of Reaching and Grasping Development." *IEEE Transactions on Autonomous Mental Development* 4 (5): 326–336.
- Scheier, Christian, Rolf Pfeifer, and Yasuo Kuniyoshi. 1998. "Embedded Neural Networks: Exploiting Constraints." *Neural Networks* 11:1551–1596.
- Schembri, Massimiliano, Marco Mirolli, and Gianluca Baldassarre. 2007. "Evolution and Learning in an Intrinsically Motivated Reinforcement Learning Robot." In *Proceedings of the European Conference on Artificial Life*. Berlin: Springer.
- Schlesinger, Matthew. 2004. "Evolving Agents as a Metaphor for the Developing Child." *Developmental Science* 7:154–168.
- Schulman, John, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. "Trust Region Policy Optimization." In *Proceedings of the 32nd International Conference on Machine Learning*, 1889–1897. PMLR.
- Schulman, John, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. "Proximal Policy Optimization Algorithms." ArXiv: 1707.06347.
- Schwefel, Hans-Paul. 1997. *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie*. Vol. 26. Basel/Stuttgart: Birkhaeuser.
- Simione, Luca, and Stefano Nolfi. 2017. "Achieving Long-Term Progress in Competitive Co-evolution." In *Proceedings of the IEEE Symposium on Computational Intelligence*. New York: IEEE.
- Simione, Luca, and Stefano Nolfi. 2019. "Long-Term Progress and Behavior Complexification in Competitive Co-evolution." ArXiv: 1909.08303.
- Sims, Karl. 1994. "Evolving 3D Morphology and Behavior by Competition." *Artificial Life* 4:28–39.
- Sinervo, Barry, and Curt M. Lively. 1996. "The Rock-Paper-Scissors Game and the Evolution of Alternative Male Strategies." *Nature* 380:240–243.
- Spencer, John P., and Esther Thelen. 2000. "Spatially Specific Changes in Infants' Muscle Coactivity as They Learn to Reach." *Infancy* 1 (3): 275–302.
- Sperati, Valerio, Vito Trianni, and Stefano Nolfi. 2011. "Self-Organised Path Formation in a Swarm of Robots." *Swarm Intelligence* 5:97–119.
- Spivey, Michael. 2007. *The Continuity of Mind*. New York: Oxford University Press.
- Stanley, Kenneth O., and Risto Miikkulainen. 2002. "Evolving Neural Networks through Augmenting Topologies." *Evolutionary Computation* 10 (2): 99–127.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Thelen, Esther, and Linda B. Smith. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Todorov, Emanuel, Tom Erez, and Yuval Tassa. 2012. "Mujoco: A Physics Engine for Model-Based Control." In *Proceedings of the 2012 IEEE/RSJ Intelligent Robots and Systems Conference*. New York: IEEE.
- Tuci, Elio, Gianluca Massera, and Stefano Nolfi. 2010. "Active Categorical Perception of Object Shapes in a Simulated Anthropomorphic Robotic Arm." *Transaction on Evolutionary Computation Journal* 14 (6): 885–899.
- Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991. *The Embodied Mind*. Cambridge, MA: MIT Press.
- von Hofsten, Claes. 1982. "Eye-Hand Coordination in the Newborn." *Developmental Psychology* 18 (3): 450–461.
- von Hofsten, Claes, and Louise Rönnqvist. 1993. "The Structuring of Neonatal Arm Movements." *Child Development* 64 (4): 1046–1057.

5 Swarm Robotics

Mary Katherine Heinrich, Mostafa Wahby, Marco Dorigo, and Heiko Hamann

5.1 Introduction

Swarm robotics is the study of how independent robots can interact as a group, giving rise to collective behaviors that a single such robot could not achieve on its own (Dorigo et al. 2014; see figure 5.1). The field can be considered an application of swarm intelligence, as defined by Bonabeau et al. (1999), and its approaches to robot control are typically based on principles of self-organization (Hamann 2018b). Swarm robotics is studied in pursuit of the oft-cited benefits that distributed or self-organized control can provide, in particular: robustness, flexibility, and scalability.

5.1.1 What Is a Swarm?

A swarm is a system of agents, whether natural or artificial, in which the characterizing behaviors occur at the group level rather than the individual level. An agent (e.g., a particle, insect, person, or robot), as defined by Russell and Norvig (2016), is “just something that acts,” and typically, it acts autonomously. Though systems of agents may show swarm behaviors that vary considerably, these behaviors are unified by their characteristic level of organization.

Swarm behaviors are not organized by a central entity that dictates instructions to individuals and likewise are not directly organized by the individuals themselves. Rather, swarm behaviors arise from the complex nonlinear dynamics of local interactions occurring in a distributed and decentralized system. Such dynamics are studied in many fields (cf. Bar-Yam 1997), being both observed in natural systems and developed in artificial ones. In nonliving systems studied in physics, self-organization can, for instance, be observed in Rayleigh-Bénard convection—wherein heating a fluid layer from below induces the formation of regular cellular patterns—or in self-organized criticality, which is seen, for example, in the power-law probability distribution of avalanche sizes. Swarm behaviors are notably widespread in biology—for instance, in social organisms. They include mobility behaviors such as flocking birds or marching locusts and spatial manipulation behaviors such as foraging in ants and honeybees or construction in termites and wasps. Artificial swarm behaviors have been studied for a broad range of tasks, including foraging (Pinciroli et al. 2012; see also the link in the additional resources section), object retrieval (Dorigo et al.



Figure 5.1

Example of a robot swarm consisting of the Kilobot. *Source:* From Rubenstein et al. 2012.

2013), and construction (Werfel et al. 2014), and recently have even been investigated for hybridization with natural systems (Hamann et al. 2017). Although real-world artificial swarms have rarely been deployed according to publicly available information, recent exceptions—in particular the NASA (2015) swarms of nanosatellites—suggest their application may become more common.

Because the behaviors that characterize a swarm occur at the group level, they can only be observed with a minimum of three agents, and many swarm behaviors will require far more. A precise definition of swarm size, as provided by Beni (2004), would be a system that is best represented as a multi-body problem, as it is respectively too large and too small for its dynamics to be well described as a few-body problem or by mean-field approximation.

5.1.2 Self-Organization and the Micro-Macro Link

Self-organization is the mechanism by which macroscale (i.e., global or systemwide) spatial and temporal structures can generate from microscale (i.e., local or peer-to-peer) interactions. In physical and biological systems, it can be observed in Rayleigh-Bénard convection, cell differentiation and embryogenesis, or pigmentation patterns in animals. In self-organizing systems, the macrostructures may generate from a combination of short-range and long-range interactions, as seen in reaction-diffusion models of biological activation-inhibition mechanisms (Meinhardt and Gierer 2000).

In swarm robotics, the link between micro- and macroscale occurs not directly but via self-organization. Actions of individual robots (i.e., microscale, or local) are typically primitive and involve a high degree of uncertainty, as they are informed only by limited knowledge and by short-range sensing and communication (see section 5.2). By contrast, collective actions of the robot swarm (i.e., macroscale, or global) are more sophisticated

and are capable of solving complex tasks. The macroscale is where development, testing, and analysis of swarm behaviors take place, while the implementation of robot controllers (i.e., executable code) occurs at the microscale. Therefore, in developing robot swarms, the desired global behaviors must be translated to local controllers, but due to the nonlinear dynamics of self-organization, this task is challenging. No generalized method has currently been developed to compile macro specifications into micro implementations, and the management of the “micro-macro link” is a key challenge in developing self-organized control (see section 5.3).

Self-organization, as defined by Bonabeau et al. (1999), functions via certain features that must be present in the system. These features are positive and negative feedback, fluctuations (i.e., random events), and multiple microscale interactions. Multiple interactions are an evident requirement, as self-organized behaviors arise from them. Positive and negative feedback are necessary to modulate deviations in the system and work in tandem to steer a robot swarm toward equilibrium or consensus. Positive feedback on its own will continuously reinforce a trend that may be based on a minor random deviation and, in all cases, will eventually surpass the desired target, creating what we might call a snowball effect or bubble. The incorporation of negative feedback is crucial to damp overshoots and tempers the impact of random deviations. Fluctuations are manageable in a swarm because of positive and negative feedback, but they are also a necessary feature, as they enable a balance between exploration and exploitation. Exploration allows a swarm to search for desired targets, while exploitation allows it to remain at those targets once they are found; a balance of these two tendencies stops a swarm from getting “stuck.” For instance, if positive feedback in a robot swarm steers it to exploit a reasonably good solution, fluctuations will be crucial for the swarm to escape that local optimum and find a better one. Likewise, if a swarm has found the best solution for a current environment, fluctuations allow it to adapt to future environmental changes by discovering that a different solution has since become superior.

5.1.3 Cognitive and Bioinspired Machine Behavior

Artificial swarms were originally heavily inspired by processes observed in biology. For instance, the dynamics governing flocks of birds, herds of mammals, and schools of fish were the inspiration behind the Reynolds (1987) model for multiagent computer graphics. Another key biological inspiration has been *stigmergy*, as seen, for instance, in ant colonies (Bonabeau et al. 1999). Stigmergy is a class of mechanisms whereby social insects do not communicate directly but rather communicate by modifying their environment in response to its current configuration, inducing nonlinear cascades of behaviors and environmental changes. In cognitive science, key perspectives such as that of Couzin (2009) consider natural stigmergy to be a cognitive mechanism. It is often considered a model of group cognition not only for insects but also for other domains such as social systems, and it can operate with many types of environmental features. In ant colonies foraging for food, for instance, the environmental cues involved in stigmergy are pheromone trails left collectively by the ants. Stigmergy has inspired the artificial swarm metaheuristic of “ant colony optimization” (Dorigo and Di Caro 1999), along with many swarm robotics approaches such as termite-inspired construction in response to the observed shape of a climbable structure (Werfel et al. 2014).

Cognitive sources of bioinspiration are common in swarm robotics and self-organized control. Slime mold, a type of amoeboid organism that spatially navigates via the self-organization of thousands of cells, does not have internal memory. It instead uses a form of spatial, external memory to steer its exploration. Slime molds have been used as models for optimization and self-organization generally. What may be considered minimal cognition in plants has in part inspired self-organized grammars such as Lindenmayer systems. The distributed steering of plant morphology in response to stimuli has inspired a “vascular morphogenesis controller” used for adaptation in robot swarm aggregation (Divband Soorati et al. 2019). In social insects, collective behaviors other than stigmergy have also been studied, such as thermoregulatory behaviors in honeybees, which have inspired “Beeclust” (Schmickl and Hamann 2011) control for robot swarm aggregation. Further inspiration may come from human social and economic systems, which are increasingly considered a form of swarm intelligence and often involve social cognition. Models used for social systems are also implemented in robots—for example, the voter model, used on online social networks, is often used for decision-making in a robot swarm (e.g., Valentini et al. 2014). Further cognitive sources of bioinspiration are discussed in sections 1.5 and 1.6.3, as emerging perspectives consider swarms a “liquid brain” class of cognition (Piñero and Solé 2019) or, alternatively, renew their consideration as a “superorganism” to which psychology models and theories can be applied (Reina et al. 2018).

5.1.4 Scalability

Changing a system’s size can cause problems. A system that is too large may have low performance due to bottlenecks, while a system that is too small may have low performance due to limited opportunities for collaboration between entities. Parallel computing defines speedup as $S = T_1/T_N$, where T_1 is the time one has to wait for the result of computation using one CPU, and T_N is the time for the same computation using N CPUs. In a presumably ideal case, one achieves linear speedups of $S = N$; a doubled system size results in doubled performance. A multirobot system is scalable if the same control algorithm can be used for both large and small numbers of robots while obtaining reasonable speedups (Hamann 2018b). Although any nontrivial multirobot scenario requires some coordination among the robots (see section 5.2), coordination can be avoided by preassigning areas of operation to each robot. This way, a multirobot scenario is effectively broken down into multiple single-robot scenarios in the form of trivial parallelization. However, if we want to make the system robust against robot failures (see section 5.1.5), then each robot should check the operation areas of other robots to see whether they accomplish their respective tasks. This requires online coordination to administrate task allocation. Disallowing collaboration between robots would also exclude the possibility of generating superlinear speedups (cf. Hamann 2018a).

Swarm robotics research targets maximal scalability—that is, the possibility of scaling to virtually any system size. The necessary requirements to achieve this are a strictly decentralized approach and limited communication. All robots exclusively use local communication and local information. Instead of point-to-point communication across the whole swarm, robots are restricted to only communicate with neighbors (“narrowcast”). If the robot density $\rho = N/A$ (number of robots N per area A) is constant, then the neighborhood size is constant,

and scaling the system is a change in the number of robots N and the provided area A . Even if these requirements are satisfied, a robot swarm may still fail to scale perfectly due to limited shared resources (e.g., the entrance to a base station) or because the required information cannot be propagated through the swarm quickly enough (e.g., by diffusion). Conversely, an advantage is that speedups of $S > N$ can potentially be achieved when robots collaborate (Hamann 2018a), for example, to cross a gap or to manipulate objects.

5.1.5 Fault Tolerance

For any engineered system, but especially in robotics, it is challenging to prepare for failures and unanticipated changes in the environment. As a simple definition, fault tolerance is a system's ability to continue functioning despite the occurrence of faults and failures. Multi-robot systems have a supposedly higher degree of fault tolerance than a single robot due merely to the system's inherent redundancy; this applies even more to robot swarms. In swarm robotics, losing one or more robots is supposed to have a limited impact on performance. Because the system is decentralized, each robot relies on local information only, and all or many robots can take over the task of another robot. The high potential for fault tolerance in robot swarms is illustrated by comparing the vulnerability of single space probe missions to the concept of swarms of nanosatellites (NASA 2015). Winfield and Nembrini (2006) have shown that the potential for fault tolerance in robot swarms has possibly been overestimated and is not necessarily an inherent feature. Partial failures of robots may be harmful, and systemwide vulnerability to faults can occur, even in robot swarms.

In a study on fault tolerance and fault detection, Christensen et al. (2009) leverage multiple equivalent units, letting them monitor each other and detect anomalies. Features that have been defined to describe robot behaviors are first determined by each robot for its neighbors and used to detect faulty behavior. In a second step, the robots collectively determine whether a robot should be considered faulty and consequentially ignored. Faulty robots decrease swarm size, such that fault tolerance requires the swarm to adapt online to changes in size. Recently, Wahby et al. (2019) have proposed a mechanism that allows robots to continuously monitor the swarm density. If a considerable change is detected, each robot adapts the parameters of its control algorithm to compensate for the changed density. In summary, there is high potential for fault tolerance in swarm robotics, but it is not inherent in all cases. Each robot is required to monitor its neighbors and relevant environmental features to detect faults or crucial changes and to adapt accordingly.

5.2 Robot-Robot Interaction

Based on the added value of automating a task with a single robot, it can seem advantageous to add another robot, and then many more. The subsequent question is whether and how the robots should interact. Allowing the robots to interact and collaborate can introduce considerable complexity to the system. One option to avoid increased complexity is the simple parallelization of tasks, with negligible communication. In a cleaning task, for example, each robot might be assigned a separate area, so that there cannot be interference between robots. One might then argue that zero interaction between robots is ideal, as this keeps the system simple. However, robot-robot interaction brings many advantageous

possibilities, such as true collaboration between robots or a performance increase that goes beyond parallelization.

In multirobot systems, different forms of robot-robot interaction can result in the emergence of collective behaviors for given tasks. These forms of interaction can be the following: direct, using explicit signaling; indirect, based on observed change in behavior or cues left in the environment; or simple physical contact. Robots interacting by physically connecting and docking to one another have been studied in reconfigurable modular robotics and in a robot swarm inspired by self-assembly in ants (Groß and Dorigo 2009). The remainder of this section describes methods of direct and indirect communication (i.e., not exclusively physical contact).

5.2.1 Direct Communication (Signals)

Robots might need to communicate their strategic decisions, progress, environmental perspective, or presence. In multirobot systems with centralized control, the robots use global communication to negotiate a strategy and assign roles. In robot swarms, by contrast, communication is constrained to be local. Therefore, infrared communication is a popular method for signaling, as it allows reliable short-range obstacle detection, distance calculation, and data communication. For example, in the Beeclust (Schmickl and Hamann 2011) control algorithm, inspired by the hive navigation behavior of young honeybees, infrared short-range obstacle detection is used for aggregation according to luminance. In the Beeclust algorithm, robots perform a random walk while turning away from obstacles and pausing when encountering another robot. A paused robot uses its luminance sensor reading e to determine the waiting period w , according to $w(e) = w_{\max} e^2 / e^2 + k$. Using infrared communication for kin recognition and communication of environmental perceptions, Wahby et al. (2019) extended the Beeclust algorithm to achieve adaptive aggregation in dynamic conditions (see figure 5.2). Other common signaling methods include short-range radio communication and visual communication via LED color. For instance, Groß et al. (2006) used blue-and-red light signals



Figure 5.2

A swarm of nine robots adapting their behavior according to detected conditions in an aggregation task. *Source:* From Wahby et al. 2019.

to influence the formation of self-assembled connection patterns for a robot swarm. Other methods such as odor and sound have also been studied. For instance, inspired by necrophoric pheromone communication that triggers corpse-removal behavior in bees, Purnama-djaja and Russell (2005) have built two mobile robots that perform a rescue task, equipped with tin oxide gas sensors. By using an odor localization algorithm, the robots can find and rescue a plastic foam artificial robot (i.e., a malfunctioning robot replica) that is releasing a chemical vapor.

5.2.2 Indirect Communication (Cues)

In indirect communication approaches, a robot in a swarm does not explicitly signal other robots or directly exchange data. Instead, the robots adjust their behavior based only on their observations of the local environment. These observations can relate to changes in the behavior of other robots or changes made to the environment (i.e., stigmergy; see section 5.1.3). Several indirect approaches have also been used to implement flocking behaviors without estimating neighbors' relative headings. For instance, Ferrante et al. (2012) defined attraction/repulsion dynamics for linear and angular velocities based only on range and bearing proximity values. Similarly, Yasuda et al. (2014) defined a topological interaction model that relies only on the proximity of local neighbors. In these approaches, the interaction is based only on the observed changes in the movement behavior of peer robots, and the robots adapt their motion accordingly.

5.2.3 Challenges of Communicating Robots

Communication is essential to allow robots to collaborate but can also be a potential bottleneck when dozens, or even hundreds, of robots need to communicate. Radio and sound communication both suffer from interference if prohibitively many senders operate in bounded areas simultaneously. Many protocols for radio communication also have further limitations and do not scale easily (e.g., Bluetooth, carrier-sense multiple access with collision avoidance [CSMA/CA]). Therefore, many swarm robotics implementations rely on other forms of communication, such as infrared with limited range (typically less than 15 cm) narrowcasting to direct neighbors.

Various hardware platforms also come with their own respective challenges. For example, aerial drone-based search and rescue missions operate in detrimental environments while requiring high bandwidth and an extensive communication range to transfer real-time footage. Many typical communication techniques are highly limited in such cases. For instance, Wi-Fi supports up to 250 m for outdoor communication, which can be a limitation in search and rescue missions that can extend to several kilometers between neighboring drones. Worldwide interoperability for microwave access network (WiMAX) technology supports a communication range of up to 30 km and is therefore a good candidate for tasks in which drones require long-range communication. The capabilities of current 4G+ mobile networks can also support the coverage and transfer rates of drones deployed at low altitudes. As a further improvement, upcoming 5G networks may provide more robust and effective connectivity for long-range communication in swarms of drones.

Underwater communication is more challenging than aerial communication because water absorbs most electromagnetic radiation except for a portion of the visible spectrum.

This visible light can furthermore travel only a few hundred meters in clear water and much shorter distances in cloudy water. Reliable acoustic modems have therefore been developed for long-range underwater communication and have been used in a swarm of autonomous underwater vehicles for communication and navigation (Behrje et al. 2018). Recently, Farr et al. (2010) have developed an optical communication method based on exchanging packets of modulated blue-green light. This method is faster and cheaper than acoustic modems but supports a shorter communication range.

5.3 Methods of Designing Robot Swarms

Designing controllers for robot swarms can be approached in the following two key ways: either with the human designer in the loop or automatically based on methods of optimization or machine learning. Both options can be challenging because of the micro-macro problem. Collective effects of many robot-robot interactions are difficult to anticipate analytically, and similarly, macroscale rewards used in automatic design cannot easily be traced back to behaviors of individual robots (see section 1.1.2).

5.3.1 Design with the Human in the Loop

The traditional approach of designing and implementing robot control algorithms is, of course, based on keeping the human in the loop; in other words, a human engineer programs the robot. In swarm robotics, often but not necessarily, control of the individual robot is kept simple because system complexity on the macroscale is supposed to emerge from robot-robot interactions. Therefore, focus has been placed primarily on simple reactive control without memory and the frameworks of behavior-based robotics. Often robot swarms have a controller based on a finite state machine. Designing a simple state-machine controller for a robot swarm is usually challenging because of the micro-macro problem (see section 5.1.2).

Even experienced robot swarm programmers need to follow an iterative trial-and-error process until the parameters of the algorithm are fine-tuned and the desired swarm behavior is achieved.

Some approaches introduce mechanisms to allow the robots to automatically adapt the parameters of a manually designed algorithm, at runtime, according to the surrounding conditions (e.g., Wahby et al. 2019). However, these approaches offer adaptive solutions tailored for task-specific scenarios and could fail in scenarios with unanticipated features. An intermediate next step before applying an automated approach is to support the human designer with models. While a trial-and-error approach uses robot simulations to estimate the result of the current algorithm design, another approach is to instead increase the level of abstraction and use a model of swarm dynamics. The objective of the modeling approach is to get generic predictions of swarm behavior for a given algorithm, rather than episodic samples from simulations. Probabilistic macroscale models are often used. The challenge is to find models that are abstract but still allow for a clear connection to the underlying control algorithm. For example, Hamann and Wörn (2008) modeled space and allowed for a mathematical connection between micro- and macroscale.

5.3.2 Evolutionary Swarm Robotics

Among automatic approaches to swarm design, artificial evolution—originally inspired by evolutionary biology—can be considered the most widespread. Evolutionary robotics (Nolfi and Floreano 2000) is a commonly followed approach outside of swarms (see chapter 4) and has been considered a framework to study generalized models of cognition (Harvey et al. 2005). The typical evolutionary swarm robotics approach is to evolve an artificial neural network controller (i.e., neuroevolution) in simulation (see link to the MultiNEAT software library in the additional resources section) in a homogeneous swarm (e.g., Baldassarre et al. 2003). Finite-state machines have also sometimes been evolved, instead of the typical neural network. A main challenge in evolutionary robotics in general, but especially in swarms, is the transfer to reality, as the evolutionary process can exploit any errors in the modeling of the experimental setup, thereby overfitting to the simulation. This “reality gap” can be addressed using the Koos et al. (2012) “transferability approach” (i.e., evaluating the evolved controllers both in simulation and in the real setup), seen, for instance, in the swarm scenario explicated in section 1.4. Online evolution (i.e., embodied evolution) is attractive for its accuracy but unattractive for its slowness, which is exacerbated in swarms. A solution to this conflict has been proposed by O’Dowd et al. (2011) via coevolution of the controller with the respective simulator. Automatic design approaches besides evolution exist, such as the modular control architecture “AutoMoDe,” where a probabilistic finite-state machine comprises a priori parametric modules wired by an optimization process (Francesca et al. 2014).

5.3.3 Neuro- and Bioinspired Automatic Design

Some inspiration sources for robot swarms have also inspired heuristics. For instance, particle swarm optimization inspired by flocking has been used in distributed versions for multirobot learning (Di Mario et al. 2015).

Artificial neural networks (ANNs)—roughly neuroinspired—have proven highly effective in many fields and have also been explored in swarm robotics (for the related topics of machine learning for robotics and neurorobotics, see chapters 3 and 9.) In a common approach, each robot in a swarm receives the same full ANN controller, evolved off-line. The “odNEAT” approach by Silva et al. (2015) extends to neuroevolution that is online and decentralized. Distributed neural networks have also been proposed. In the approach of Otte (2018), each robot holds a slice of neurons in a swarm-wide ANN, enabled by parallel neural network training. In an alternative neuroinspired approach, Mathews et al. (2017) have developed “mergeable nervous systems,” where attached robots can flexibly fuse their distributed control systems into a shared adaptive network.

5.4 Indoor and Outdoor Applications of Robot Swarms

Swarm robotics research often focuses on fundamental models and design approaches, supported by experiments in laboratory environments. Although basic characteristics of robot swarms, such as scalability, would evidently have an impact on applications, specific applied scenarios have rarely been studied directly. Some approaches have indirectly studied a specific industrial or field task despite conducting only laboratory experiments.

For an industrial task, a laboratory approach can use a stand-in robot to replicate the key sensing and actuation capabilities of a patented industrial robot and then use the laboratory stand-in to study self-organized control (e.g., reconfigurable fiber deployment in manufacturing; Eschke et al. 2019). Swarm robotics approaches in laboratory environments have also proposed solutions to field tasks—such as the problem of impassable step height in disaster relief—for example, by distributed construction of amorphous ramps (Napp and Nagpal 2014). In another approach, elements from the field can be brought into laboratory environments for experiments, as seen in biohybrid robotics research with plants (Wahby et al. 2018).

It is recently becoming more common for swarm robotics research to conduct field experiments. The “subCULTron” EU project (Thenius et al. 2016) is testing its swarm of underwater robots for marine monitoring in a lagoon environment in Venice, Italy (see in-process field photos printed in Hamann [2018b]). Another project, “SAGA,” develops a swarm of quadrotor UAVs for field monitoring and mapping of agricultural conditions—for instance, with weed detection (Albani et al. 2017).

5.4.1 Example Outdoor Scenario

In order to provide a didactic example of an indoor or outdoor application, we give a detailed walk-through of an approach by Duarte et al. (2016) because it is the first published instance of real field experiments with a robot swarm. Duarte et al. (2016) use ten autonomous aquatic surface vehicles and test them in a shallow open-water environment in Lisbon, Portugal. The robots are differential drive boat vehicles that use inexpensive and accessible off-the-shelf components. They are 60 cm at their longest dimension, are capable of up to 1.7 m/s linear speed and 90°/s rotational speed, and comprise components costing roughly three hundred euros per robot. Each boat robot is equipped for decentralized communication with other robots via a wireless ad hoc network for UDP (User Datagram Protocol) broadcasting and is equipped with GPS and a compass. The controllers output linear and rotational speeds, which are used to calculate motor speeds based on the real robot dynamics (affected by friction and inertia in water). The controller inputs are three values representing locations in the environment, calculated from GPS and compass readings of the robot and its neighbors, as communicated over the wireless network.

Using these robots and controllers, Duarte et al. (2016) have conducted simulated and real field experiments for four different tasks that require coordination between robots. The robot controllers are evolved in simulation,¹ then transferred to real field experiments in open water using the transferability approach of evolutionary robotics (see section 5.3.2). In the first task, homing, the swarm collectively moves to a target in the environment while avoiding collisions. During evolution, controllers are rewarded for minimizing distance d to the target; specifically, the average value of $\Delta d/d_{t=0}$ for each robot at each time step, multiplied by coefficient S to penalize controllers when robots get less than 3 m apart. The second and third tasks are dispersion and aggregation, in which the robots should either spread out over a large area without losing contact with neighbors or should move toward each other to form clusters after starting from a spread-out configuration. The fourth task is area monitoring, in which the robots should move around to collectively give continual coverage to a defined and limited area. The four behaviors are combined into a single “multicontroller” mission in the field, which was not previously evolved for

or tested in simulation. The researchers equip the robots with temperature sensors for this mission and select the highest-performing controllers from each respective task. The four controllers are triggered sequentially in the swarm, successfully completing an application-oriented mission of sampling water temperature. Within this mission, the robot swarm moves in a close group from the starting point to the target area, disperses and monitors the full area, then aggregates back into a close group and returns to the initial starting point.

5.5 Swarm Cognition and Psychology

As introduced in section 5.1.3, collective cognition is found in many natural swarms and is a target in engineering artificial ones. An established perspective on natural swarms is that their collective behaviors bear commonalities with neural mechanisms and therefore should be studied in the same framework of cognitive science (Couzin 2009; Trianni et al. 2011). Another perspective holds that swarms should be studied as an independent class of cognition, forming what can be considered “liquid brains” (Piñero and Solé 2019).

Processes of collective cognition that are investigated in swarms include collective perception (Schmickl et al. 2006), collective memory (Couzin et al. 2002), collective learning (Montes de Oca and Stützle 2008), and collective decision-making (see section 5.6). Cognitive processes observed in simple organisms that rely on decentralization, such as ants, have commonly inspired swarm robotics. Examples inspired by more complex organisms, or by coordination that is not strictly decentralized, are far more rare. However, there are a few examples. Regarding more complex organisms with higher-order cognition and centralized nervous systems, there has been inspiration from neuroscience (e.g., in automatic design methods for swarms) and human psychology (e.g., in natural swarms that can be considered superorganisms). Regarding coordination that is not strictly decentralized, species with hierarchical social structures (e.g., baboons) display coordination strategies that may, speculatively, be relevant to multirobot control. It has also been proposed that simpler social animals such as schooling fish, often considered to exclusively use peer-to-peer communication for movement, may sometimes use hierarchical social structures with temporary leaders for fast predator response (Ioannou 2017). We therefore look to neuroscience and human psychology—in addition to models of complex social structures such as those seen in online social networks or hierarchical animal groups—for key theories that may have potential for useful application in a robot swarm.

Key theories from psychology and neuroscience have thus far been implemented in models of swarm behavior in a few seminal works on collective decision-making, described in detail in section 5.6.3. Implementations of such theories have not occurred in models of swarm perception, memory, or learning. We therefore describe existing swarm robotics examples related to these aspects of cognition and review some of the key psychology and neuroscience theories that are potentially relevant to distributed and decentralized robot cognition. As our aim is to follow a natural inspiration source only insofar as is useful for the engineering task at hand, we present theories based on their potential relevance to robot control, without taking a stance on the positions of those theories within their originating disciplines.

5.5.1 Collective Perception and Attention

In existing strategies for collective perception in a robot swarm, peers trade information capturing their individual perceptions with their local neighbors, progressively building consensus about the perceived environment. For instance, they signal votes or hypotheses about perceived features (Valentini et al. 2014) or share “trophallaxis-inspired” cues about implicit elapsed time since they last reached a target (Schmickl et al. 2006). It is typically held that natural swarms similarly use distributed strategies for perception. However, it is sometimes conversely held that in some social animals, such as fish, the improved predator perception of larger groups may result simply from a pooled visual field and the temporary leadership of a faster-moving individual (Ioannou 2017), without any peer-to-peer communication about perception.

Established human psychology laws for stimuli-response mechanisms have been shown to be relevant to collective decision-making—for instance, in terms of the speed-accuracy trade-off in swarms—and may also relate to collective perception. In disciplines such as human-computer interaction, motor speed-accuracy trade-offs have been well described by the psychology principle of Fitts’s law, proposed by Paul Fitts in 1954, wherein the size and distance of a target predict movement patterns toward it. Though established as a motor law, it has been shown to hold for agents’ perception of action (Grosjean et al. 2007), an important aspect of robot-robot collaboration in swarms. As another example, psychology has established a relationship between attention levels and the exploration-exploitation trade-off in foraging (Van den Driessche et al. 2019), a task often studied in swarm robotics.

In biology, sensorimotor processes are key to perception, especially in coordination between individuals. Santana and Correia (2010) propose that, by considering attention in isolation from subsequent motor system processes, biological neural mechanisms might inspire approaches to swarm perception. For example, mechanisms governing selective attention could be transferred to robot swarms to establish a relationship between attention behaviors and predictions or motivations.

5.5.2 External and Collective Memory

Behavioral science has proposed a variety of group memory concepts in natural swarms, such as the “joint memory” proposition of Thierry et al. (1995), including, first, a collective type in which memories of individuals are coupled and, second, an external type in which memory refers to the environment. External memory might be saved in modifications to the environment, as in stigmergy, or may simply comprise references to landmarks in the environment (e.g., to facilitate novel actions rather than the repetition of remembered actions). In artificial swarms, a simple approach is to equip agents with local memory of their own history to enhance performance when interacting with the environment. Another approach, which can be applied to foraging in robot swarms, is the use of a maplike representation of terrain features, which may be predetermined or built adaptively (Kumar and Sahin 2003). The most common approach in robot swarms is evidently the external memory approach of stigmergy, which can also be combined with other types of memory, such as short-term memory of individual history. Short-term memory in a swarm is discussed further below, in relation to a natural swarm being considered a superorganism.

Theoretical biology notably provides simulation-based evidence of collective memory in swarms, demonstrating that the history of swarm structure has an impact on the different collective behaviors that might arise from identical individual behaviors (Couzin et al. 2002). In honeybees, Beekman (2005) has experimentally demonstrated individual memories of past stimuli that may affect subsequent interactions and collective behavior, as agents triggered by others to revisit a site that is still remembered will be more efficient (e.g., by avoiding unsuccessful route attempts).

In the coupled-memory type proposed by Thierry et al. (1995), individuals manage their own memory of an opinion or piece of information and communicate that individual memory to others—for instance, in honeybees, each knows only a portion of information about an environment. In existing robot swarms, there are typically no subgroups of spatial memory distribution across a swarm (i.e., the opinions held by individuals vary, but the topic on which they have opinions is homogeneous). However, the role that an individual plays in information processing in a swarm may be influenced by its spatial position. It is notable that distributed memory in the brain is heavily differentiated according to spatial distribution, but the physical connections present in biological neural circuits may limit them as a direct inspiration source for robot swarms. In social insects, differentiated memory subgroups have been shown to arise, specifically, when a small group of individuals becomes short-term specialists for a repeated, temporary task (Diez et al. 2011).

5.5.3 Social and Collective Learning

Social learning, or collective learning, refers to the process of behavior development via observation and imitation of neighbors (Rendell et al. 2010). A common mechanism in swarms that may be considered a simple form of social learning is the disproportionately large influence that a few informed individuals have on the behavior of a group. The proportion of informed agents needed to maintain accuracy has even been shown to decrease with increasing group size (cf. Couzin 2009). Procedures to reach consensus in collective decision-making (addressed as its own aspect of cognition in section 5.6) have also been considered a type of social learning in animal groups, in cases in which agents are selective about the neighbors they imitate (Rendell et al. 2010). This selectivity has roughly inspired a social-learning approach in artificial swarms, where a large group reaches consensus more quickly by incrementally adding agents to an initially small decision-making subgroup (Montes de Oca and Stützle 2008). In another approach, artificial agents follow instructions from a leader and use these downstream instructions to indirectly learn the respective task so they can collectively reconstruct it if the leader is lost (Karydis et al. 2016).

In social animals, associative learning in an individual has been frequently studied, establishing a direct link between individual preferences and actions. However, Kao et al. (2014) contend that the majority of the animals studied in lab conditions will naturally exist in social groups, where collective learning will break the established relationship between preference and action in associative learning. The influence of collective learning on associative learning in animals has yet to be studied directly (Kao et al. 2014), although related established studies on honeybees have examined both associative learning by cues and social learning by the well-known mechanism of dance communication. The effect of

agent individuality (i.e., behavioral heterogeneity) on natural swarm dynamics has been studied, possibly bringing implications for robot swarms (Saffre et al. 2018).

Burini et al. (2016) have proposed a unified formulation of collective-learning dynamics using kinetic theory, including learning of abilities and of social messages. Their formulation presumes heterogeneity in the group (i.e., the “population-thinking” approach)—in a robot swarm, such heterogeneity could potentially be characterized as deviations in behaviors or opinions during progression toward consensus. Approaches to opinion consensus in robot swarms have been studied in collective decision-making.

5.6 Collective Decision-Making in Robots

Collective decision-making is the key mechanism of swarm cognition. A robot swarm can only act as a whole by ensuring consensus or vast majorities for certain coordinated actions. Achieving such consensus and coordination in a swarm, particularly in unknown or dynamic environments, requires swarm-wide sensing, information processing, and action selection.

5.6.1 Swarm Autonomy and Swarm Awareness

Following the agent models of Russell and Norvig (2016), the autonomy of an agent originates from its ability to make informed decisions. Similarly, a robot swarm can only be autonomous and self-governing on a macroscale if the swarm as a whole is capable of making informed decisions. This requires a form of collective decision-making that ensures the collection of relevant information, collective processing of that information, and a subsequent swarm-wide decision of what to do next. In addition, the swarm needs to reach awareness that a decision is necessary and that a consensus or large majority has been achieved such that the decision-making process concludes (Hamann 2018b). As pointed out, for example, by Khaluf et al. (2019), this corresponds to common subdivisions of human decision-making, such as identifying the problem, obtaining information (identification of options and their quality), and evaluating it.

In swarm robotics, and also in opinion dynamics and related fields, some aspects of swarm awareness are often ignored (Khaluf et al. 2019). Experiments often isolate one aspect, for instance, by starting immediately with the collective decision-making process before being stopped by an external observer once a sufficient majority is reached. The challenge of extending beyond experiments of isolated aspects will be crucial to achieving full swarm awareness. For full swarm awareness, each robot needs to be sensitive to changes in the environment or in the (signaled) states of its neighbors. If the swarm in a critical situation does not collectively detect that a collective decision is required, then the swarm may split, crash, or otherwise fail at its task. Similarly, to ensure the decision-making process ends, each robot needs to estimate when to stop switching opinions. As each robot relies on local information only, this estimation is necessarily probabilistic. This can be implemented as each robot voting for ending the collective decision-making process, which consequently means that we are embedding another decision-making process into the system. This can be even more challenging when the swarm has to adapt to environmental conditions and adaptively balance the speed versus accuracy of its decision-making process. So collective decisions make a swarm autonomous on a macroscale but

also require sophisticated forms of information diffusion, gossiping, and sharing of internal states to create swarm awareness of globally required swarm actions.

5.6.2 Methods of Collective Decision-Making

Methods of implementing a complete collective decision-making system include all of the following: starting and ending the process, exploring options, disseminating knowledge about options, processing that information in an individual robot, and ensuring the swarm decides accurately and quickly (Hamann 2018b). These parts are all complex and cannot be discussed in full detail here. Even more complexity would be added when considering modeling techniques that deal with underlying dynamic networks and, for instance, try to predict expected convergence times. Instead, this section focuses on the decision-making mechanism of an individual robot and the impact of different algorithm choices by looking at two simple techniques. Assuming that a robot operates iteratively on three phases (explore, disseminate, and switch opinion), here we focus only on the opinion-switch phase. Take, for instance, a robot that collects messages from its neighbors that include merely whether they are in favor of option red or option blue. Then it is reasonable to count red and blue messages, to determine the majority, and to switch to the majority opinion (or keep it if the robot already has that opinion).

This straightforward approach is called *majority rule* (e.g., Valentini et al. 2015)—that is, in a decision between two opinions, the opinion of robot r_i switches if it does not match the majority opinion in $[r_i, r_{i+n}]$, where n is the number of robots in the neighborhood. Example code for majority rule can be found in the PyCX repository (Sayama 2013). If each robot follows this simple rule, then the expectation may be that the swarm will converge on a consensus, given enough time. In general, this is true, but it can be complicated by noise or by inhomogeneously distributed robots in space (Valentini et al. 2015). A second straightforward approach is the *voter model* (e.g., Valentini et al. 2014)—that is, the opinion of robot r_i switches to a uniformly randomly selected opinion from the robots in the neighborhood, $[r_{i+1}, r_{i+n}]$. Example code for a voter model can also be found in the PyCX repository (Sayama 2013). Although it may seem counterintuitive, the voter model is a useful option for a decision-making mechanism. In decision-making, and also in collective decision-making, there exists a speed-accuracy trade-off (also mentioned at the end of section 5.6.3). This trade-off means that a decision-making process can be either fast or accurate but not both at the same time. Whether a given decision mechanism is better than another will always depend on the requirements of a given application scenario. In general, the majority rule is fast but relatively inaccurate, while the voter model is accurate but slow. There is no simple description that can provide an intuitive understanding of that finding except that the voter model tends to be more forgiving to local temporary deviations, while majority rule tends to exploit the current local system state. A better understanding requires deeper study of the different modeling techniques of collective decision-making.

5.6.3 Psychology of the Robot-Swarm Superorganism

One demonstrated approach to modeling collective decision-making in a swarm is to take inspiration from established mechanisms in human psychology and apply them to the whole swarm as if it were one organism. The group cognition and organization seen in

natural swarms has sometimes prompted their biological characterization as *superorganisms* (cf. Wilson and Sober 1989). In a superorganism, such as a honeybee colony, natural selection might operate according to the survival of the colony as a unit, evolving a tightly interdependent group and establishing a higher class of biological organization. This tighter interdependence can be seen in social apoptosis in honeybees, where colony immunity is supported by the increased infection susceptibility of individual sacrificial bees (Page et al. 2016). The superorganism concept can look similar to the established *group selection* mechanism in evolution of cooperation but has also been proposed as distinct. Without commenting on evolutionary biology, here we refer to the superorganism as a useful analogy concerning natural swarms, and potentially robot swarms. Natural swarms have been shown to perform typical organism-level functions at the level of the group, for instance, by a “common stomach” to regulate foraging (Schmickl and Karsai 2016) or neurologically by governing speed-accuracy trade-offs similarly to the brains of individual animals (Sasaki and Pratt 2018). Collective decision-making in colonies responding to stimuli has notably been shown to follow certain established laws of human psychology (Pais et al. 2013; Reina et al. 2018), a generality that may extend to robot swarms.

The signals and cues of robot swarm communication described in section 5.2 are also seen in natural swarms, as stimuli that might be shaped differently by selection—in evolutionary biology, signals are stimuli formed for the express purpose of communication, while cues are stimuli that may trigger responses but have not necessarily developed for that function. As shown by Reina et al. (2018), although stronger signals are known to lead to faster decisions and avoidance of deadlocks, they may also lead to negative performance effects. In simulated honeybee colonies, Reina et al. (2018) demonstrated that when measuring signaling by signal-to-noise ratio, increased signaling worsens the group ability to differentiate between similar stimuli. This is reminiscent of the well-known exploration-exploitation trade-off in swarm robotics. As described in section 1.1.2, exploration and exploitation are necessary mechanisms in self-organization. Achieving the optimal balance between exploiting already known solutions and exploring to find new (possibly better) solutions cannot be solved generically as it depends on the respective task and environment. For example, in a bistable regime where a robot swarm should select the best site but finds two equally good sites, the main challenge for the exploration-exploitation trade-off is to break symmetry effectively (Hamann et al. 2012).

Established psychology laws may govern the dynamics known to be present in natural swarms; for instance, Reina et al. (2018) demonstrated that the exploration-exploitation trade-off in a honeybee colony may be governed by Weber’s law on the perception of external stimuli, proposed by Gustav Fechner in 1858. This law describes differential sensitivity as $dp = dS/S$ —that is, the perceived change in stimulus dS is proportional to the initial stimulus S . In natural swarms, Pais et al. (2013) and Reina et al. (2018) have shown that Weber’s law holds in honeybee colonies choosing between sites in a bistable regime. In swarms that maintain relative spatial distributions, such as flocking birds, Perna et al. (2019) have shown that a simple antidiffusion mechanism based on Weber’s law is alone sufficient to achieve stability, compared to the two or more separate mechanisms needed to balance one another in the classic Reynolds (1987) approach. Similarly, natural flight patterns observed in honeybee colonies have been shown by Reynolds et al. (2013) to be achievable by odometry following Weber’s law.

Another established psychology law—Hick’s law, proposed by William Hick in 1952 and Robert Hyman in 1953—describes a concept termed rate of information gain, holding that reaction time rises linearly with the degree of uncertainty. That is, reaction time $RT = kH$, where H represents the amount of information that must be processed in a given decision. In the case of equally likely alternatives, Hick’s law states that $H = \log_2(n + 1)$, such that H is a logarithmic function of the number n of stimulus-response alternatives. Reina et al. (2018) have found that in a honeybee colony superorganism making a best-of- n decision, reaction time RT rises with the number of alternatives as in Hick’s law but rises exponentially, proposing that this may be due to nonlinearities in the swarm. Another established psychology law that may fit this phenomenon is the Cooney and Troyer (1994) approach that integrates interference susceptibility into a model of reaction time. Alternatively, Reina et al. (2018) have proposed a new model of reaction time RT in a honeybee superorganism: $RT = \alpha \bar{v}^{-\beta} e^{\mu n}$, where \bar{v} is the mean quality or likelihood of the n available options, and α , β , and μ are constants. This new model by Reina et al. (2018) combines Hick’s law with the Pieron law, proposed by Henri Pieron in 1913, wherein reaction time decreases with increasing intensity of stimulus as a power law.

In their implementation of Hick’s law, Reina et al. (2018) have found a trade-off in signal-to-noise ratio in a best-of- n decision, in which increased signaling improves speed but weakens selection quality, fitting with the established speed-accuracy tradeoff seen not only in robot swarm decision-making and in natural decision-making but across many aspects of information processing. A variety of factors demonstrably affect speed and accuracy in decision-making and can potentially have an impact on their trade-off in engineering robot swarms. In animal populations, the speed-accuracy trade-off during selection is proposed to result in a heterogeneous behavioral tendency to be fast or slow, as both strategies may perform similarly due to a related risk-reward trade-off (Sih and Del Giudice 2012). In natural swarms, Pais et al. (2013) have shown that in a honeybee colony where binary alternatives are distinguishable, as defined by Weber’s law, the speed-accuracy trade-off is dependent on cross-inhibition strength (a mechanism observed both in honeybee colonies and in complex brains). In individual human decision-making, when accuracy itself displays a trade-off between true and false positives, a collective approach has been shown to invert that trade-off by both increasing true positives and decreasing false ones (Wolf et al. 2013). Reina et al. (2018) have noted that accuracy in a natural swarm is dependent on the ratio between the time spent exploring versus signaling, reminiscent of the exploration-exploitation trade-off.

5.7 Conclusion

Swarm robotics was initially inspired by behaviors observed in biology, and new advances in artificial swarms continue to be interdependent with those of natural swarms, especially in the study of swarm cognition. Bioinspired and neuroinspired approaches have been used to develop robot swarm models and behaviors—such as the cognitive mechanism of stigmergy—and have influenced popular automatic design methods for swarm controllers, such as neuroevolution. Swarm robotics uses these approaches to address challenges in, for instance, direct and indirect communication, management of the “micro-macro link,”

swarm autonomy, and swarm cognition, and is moving toward applications in the field. Swarm cognition has been studied in terms of collective perception, collective memory, collective learning, and collective decision-making and, in some cases, takes inspiration from human psychology and cognitive sciences. These disciplines may provide swarm robotics with new and useful inspiration sources if measurably novel and not a reformulation of an existing approach, and if effective for the respective engineering task.

Additional Reading and Resources

- The classical introduction to swarm intelligence: Bonabeau, Eric, Marco Dorigo, and Guy Theraulaz. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford: Oxford University Press.
- A recent, comprehensive overview of swarm robotics, with detailed presentations of methods and example scenarios for the design of large-scale robot swarms: Hamann, Heiko. 2018. *Swarm Robotics: A Formal Approach*. Berlin: Springer.
- A recent perspective of swarm robotics and its future: Dorigo, Marco, Guy Theraulaz, and Vito Trianni. 2021. “Swarm Robotics: Past, Present, and Future.” *Proceedings of the IEEE* 109 (7): 1152–1165.
- A brief summary of swarm robotics’s origins, application domains, and current research issues: Dorigo, M., G. Theraulaz, and V. Trianni. 2021. “Swarm Robotics: Past, Present, and Future.” *Proceedings of the IEEE* 109 (7): 1152–1165. <https://doi.org/10.1109/JPROC.2021.3072740>.
- Software for swarm foraging, in the repository of the ARGoS simulator (Pinciroli et al. 2012): <https://github.com/ilpincy/argos3-examples>.
- MultiNEAT software library for the evolution of neural networks: <http://www.multineat.com>.
- scikit-learn software library for machine learning, including for training neural networks: <https://scikit-learn.org/stable/>.
- Software for majority-rule simulations: <https://github.com/hsayama/PyCX>.

Acknowledgments

This work has been partially supported by the Program of Concerted Research Actions (ARC) of the Université libre de Bruxelles. Marco Dorigo and Mary Katherine Heinrich acknowledge support from the Belgian F.R.S.-FNRS, of which they are a research director and a postdoctoral researcher, respectively.

Note

1. Evolution in Duarte et al. (2016) was conducted in JBotEvolver, available at <https://github.com/BioMachinesLab/jbotevolver>.

References

- Albani, Dario, Joris Ijsselmuiden, Ramon Haken, and Vito Trianni. 2017. "Monitoring and Mapping with Robot Swarms for Agricultural Applications." In *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance*, 1–6. New York: IEEE.
- Baldassarre, Gianluca, Stefano Nolfi, and Domenico Parisi. 2003. "Evolving Mobile Robots Able to Display Collective Behaviors." *Artificial Life* 9 (3): 255–267.
- Bar-Yam, Yaneer. 1997. *Dynamics of Complex Systems*. Boca Raton: CRC Press.
- Beekman, Madeleine. 2005. "How Long Will Honey Bees (*Apis mellifera* L.) Be Stimulated by Scent to Revisit Past-Profitable Forage Sites?" *Journal of Comparative Physiology A* 191 (12): 1115–1120.
- Behrje, Ulrich, Cedric Isokeit, Benjamin Meyer, and Erik Maehle. 2018. "A Robust Acoustic-Based Communication Principle for the Navigation of an Underwater Robot Swarm." In *2018 OCEANS—MTS/IEEE Kobe Techno-Oceans (OTO)*, 1–5. New York: IEEE.
- Beni, Gerardo. 2004. "From Swarm Intelligence to Swarm Robotics." In *Swarm Robotics: SAB 2004 International Workshop, Santa Monica, CA, USA, July 17, 2004, Revised Selected Papers*, edited by Erol Şahin and William M. Spears, 1–9. Berlin: Springer.
- Bonabeau, Eric, Marco Dorigo, and Guy Theraulaz. 1999. *Swarm Intelligence: from Natural to Artificial Systems*. Oxford: Oxford University Press.
- Burini, Diletta, Silvana De Lillo, and Livio Gibelli. 2016. "Collective Learning Modeling Based on the Kinetic Theory of Active Particles." *Physics of Life Reviews*, no. 16, 123–139.
- Christensen, Anders Lyhne, Rehan O'Grady, and Marco Dorigo. 2009. "From Fireflies to Fault-Tolerant Swarms of Robots." *IEEE Transactions on Evolutionary Computation* 13 (4): 754–766.
- Cooney, John B., and Rod Troyer. 1994. "A Dynamic Model of Reaction Time in a Short-Term Memory Task." *Journal of Experimental Child Psychology* 58 (2): 200–226.
- Couzin, Iain D. 2009. "Collective Cognition in Animal Groups." *Trends in Cognitive Sciences* 13 (1): 36–43.
- Couzin, Iain D., Jens Krause, Richard James, Graeme D. Ruxton, and Nigel R. Franks. 2002. "Collective Memory and Spatial Sorting in Animal Groups." *Journal of Theoretical Biology* 218 (1): 1–12.
- Diez, Lise, Jean-Louis Deneubourg, Lucie Hoesbeke, and Claire Detrain. 2011. "Orientation in Corpse-Carrying Ants: Memory or Chemical Cues?" *Animal Behavior* 81 (6): 1171–1176.
- Di Mario, Ezequiel, Inaki Navarro, and Alcherio Martinoli. 2015. "A Distributed Noise-Resistant Particle Swarm Optimization Algorithm for High-Dimensional Multi-robot Learning." In *International Conference on Robotics and Automation*, 5970–5976. New York: IEEE.
- Divband Soorati, Mohammad, Mary Katherine Heinrich, Javad Ghofrani, Payam Zahadat, and Heiko Hamann. 2019. "Photomorphogenesis for Robot Self-Assembly: Adaptivity, Collective Decision-Making, and Self-Repair." *Bioinspiration and Biomimetics* 14 (5): 056006.
- Dorigo, Marco, Mauro Birattari, and Manuele Brambilla. 2014. "Swarm Robotics." *Scholarpedia* 9 (1): 1463.
- Dorigo, Marco, and Gianni Di Caro. 1999. "Ant Colony Optimization: A New Metaheuristic." In *Congress on Evolutionary Computation*, 1470–1477. New York: IEEE.
- Dorigo, Marco, Dario Floreano, Luca Maria Gambardella, Francesco Mondada, Stefano Nolfi, Tarek Baaboura, Mauro Birattari, et al. 2013. "Swarmanoid: A Novel Concept for the Study of Heterogeneous Robotic Swarms." *IEEE Robotics and Automation Magazine* 20 (4): 60–71.
- Duarte, Miguel, Vasco Costa, Jorge Gomes, Tiago Rodrigues, Fernando Silva, Sancho Moura Oliveira, and Anders Lyhne Christensen. 2016. "Evolution of Collective Behaviors for a Real Swarm of Aquatic Surface Robots." *PLoS One* 11 (3): e0151834.
- Eschke, Catriona, Mary Katherine Heinrich, Mostafa Wahby, and Heiko Hamann. 2019. "Self-Organized Adaptive Paths in Multi-robot Manufacturing: Reconfigurable and Pattern-Independent Fibre Deployment." In *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4086–4091. New York: IEEE.
- Farr, N., A. Bowen, J. Ware, C. Pontbriand, and M. Tivey. 2010. "An Integrated, Underwater Optical/Acoustic Communications System." In *OCEANS'10 IEEE SYDNEY*, 1–6. New York: IEEE.
- Ferrante, Eliseo, Ali Emre Turgut, Cristián Huepe, Alessandro Stranieri, Carlo Pinciroli, and Marco Dorigo. 2012. "Self-Organized Flocking with a Mobile Robot Swarm: A Novel Motion Control Method." *Adaptive Behavior* 20 (6): 460–477.
- Francesca, Gianpiero, Manuele Brambilla, Arne Brutschy, Vito Trianni, and Mauro Birattari. 2014. "AutoMoDe: A Novel Approach to the Automatic Design of Control Software for Robot Swarms." *Swarm Intelligence* 8 (2): 89–112.

- Grosjean, Marc, Maggie Shiffar, and Günther Knoblich. 2007. "Fitts's Law Holds for Action Perception." *Psychological Science* 18 (2): 95–99.
- Groß, Roderich, Michael Bonani, Francesco Mondada, and Marco Dorigo. 2006. "Autonomous Self-Assembly in Swarm-Bots." *IEEE Transactions on Robotics* 22 (6): 1115–1130.
- Groß, Roderich, and Marco Dorigo. 2009. "Towards Group Transport by Swarms of Robots." *International Journal of Bio-inspired Computation* 1 (1/2): 1–13.
- Hamann, Heiko. 2018a. "Superlinear Scalability in Parallel Computing and Multi-robot Systems: Shared Resources, Collaboration, and Network Topology." In *Architecture of Computing Systems*, 31–42. Berlin: Springer.
- Hamann, Heiko. 2018b. *Swarm Robotics: A Formal Approach*. Berlin: Springer.
- Hamann, Heiko, Thomas Schmickl, Heinz Wörn, and Karl Crailsheim. 2012. "Analysis of Emergent Symmetry Breaking in Collective Decision Making." *Neural Computing and Applications* 21 (2): 207–218.
- Hamann, Heiko, Mohammad Divband Soorati, Mary Katherine Heinrich, Daniel Nicolas Hofstadler, Igor Kuksin, Frank Veenstra, Mostafa Wahby, et al. 2017. "Flora Robotica—an Architectural System Combining Living Natural Plants and Distributed Robots." ArXiv preprint: 1709.04291.
- Hamann, Heiko, and Heinz Wörn. 2008. "A Framework of Space-Time Continuous Models for Algorithm Design in Swarm Robotics." *Swarm Intelligence* 2 (2–4): 209–239.
- Harvey, Inman, Ezequiel Di Paolo, Rachel Wood, Matt Quinn, and Elio Tuci. 2005. "Evolutionary Robotics: A New Scientific Tool for Studying Cognition." *Artificial Life* 11 (1–2): 79–98.
- Ioannou, Christos C. 2017. "Swarm Intelligence in Fish? The Difficulty in Demonstrating Distributed and Self-Organised Collective Intelligence in (Some) Animal Groups." *Behavioral Processes* 141 (2): 141–151.
- Kao, Albert B., Noam Miller, Colin Torney, Andrew Hartnett, and Iain D. Couzin. 2014. "Collective Learning and Optimal Consensus Decisions in Social Animal Groups." *PLoS Computational Biology* 10 (8): e1003762.
- Karydis, Konstantinos, Prasanna Kannappan, Herbert G. Tanner, Adam Jardine, and Jeffrey Heinz. 2016. "Resilience through Learning in Multi-agent Cyber-Physical Systems." *Frontiers in Robotics and AI* 3:36.
- Khaluf, Yara, Pieter Simoens, and Heiko Hamann. 2019. "The Neglected Pieces of Designing Collective Decision-Making Processes." *Frontiers in Robotics and AI* 6:16.
- Koos, Sylvain, Jean-Baptiste Mouret, and Stéphane Doncieux. 2012. "The Transferability Approach: Crossing the Reality Gap in Evolutionary Robotics." *IEEE Transactions on Evolutionary Computation* 17 (1): 122–145.
- Kumar, Vignesh, and Ferat Sahin. 2003. "Cognitive Maps in Swarm Robots for the Mine Detection Application." In *2003 IEEE International Conference on Systems, Man and Cybernetics*, 3364–3369. New York: IEEE.
- Mathews, Nithin, Anders Lyhne Christensen, Rehan O'Grady, Francesco Mondada, and Marco Dorigo. 2017. "Mergeable Nervous Systems for Robots." *Nature Communications* 8 (1): 1–7.
- Meinhardt, Hans, and Alfred Gierer. 2000. "Pattern Formation by Local Self-Activation and Lateral Inhibition." *Bioessays* 22 (8): 753–760.
- Montes de Oca, Marco A., and Thomas Stützle. 2008. "Towards Incremental Social Learning in Optimization and Multiagent Systems." In *Proceedings of the 10th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '08)*, 1939–1944. New York: Association for Computing Machinery.
- Napp, Nils, and Radhika Nagpal. 2014. "Distributed Amorphous Ramp Construction in Unstructured Environments." *Robotica* 32 (2): 279–290.
- NASA. *Network and Operation Demonstration Satellite NASA*. 2015. Accessed July 2020. https://www.nasa.gov/mission_pages/station/research/experiments/explorer/Investigation.html?id=1601.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics*. Cambridge, MA: MIT Press.
- O'Dowd, Paul J., Alan F. T. Winfield, and Matthew Studley. 2011. "The Distributed Co-evolution of an Embodied Simulator and Controller for Swarm Robot Behaviors." *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4995–5000. New York: IEEE.
- Otte, Michael. 2018. "An Emergent Group Mind across a Swarm of Robots: Collective Cognition and Distributed Sensing via a Shared Wireless Neural Network." *International Journal of Robotics Research* 37 (9): 1017–1061.
- Page, Paul, Zheguang Lin, Ninat Buawangpong, Huoqing Zheng, Fuliang Hu, Peter Neumann, Panuwan Chantawanakul, and Vincent Dietemann. 2016. "Social Apoptosis in Honey Bee Superorganisms." *Scientific Reports* 6:27210.
- Pais, Darren, Patrick M. Hogan, Thomas Schlegel, Nigel R. Franks, Naomi E. Leonard, and James A. R. Marshall. 2013. "A Mechanism for Value-Sensitive Decision-Making." *PLoS One* 8 (9): e73216.
- Perna, Andrea, Giulio Facchini, and Jean-Louis Deneubourg. 2019. "Weber's Law-Based Perception and the Stability of Animal Groups." *Journal of the Royal Society Interface* 16 (154): 20190212.
- Pinciroli, Carlo, Vito Trianni, Rehan O'Grady, Giovanni Pini, Arne Brutschy, Manuele Brambilla, Nithin Mathews, et al. 2012. "ARGoS: A Modular, Parallel, Multi-engine Simulator for Multi-robot Systems." *Swarm Intelligence* 6 (4): 271–295.

- Piñero, Jordi, and Ricard Solé. 2019. “Statistical Physics of Liquid Brains.” *Philosophical Transactions of the Royal Society B* 374 (1774): 20180376.
- Purnamadajaja, Anies Hannawati, and R. Andrew Russell. 2005. “Pheromone Communication in a Robot Swarm: Necrophoric Bee Behavior and Its Replication.” *Robotica* 23 (6): 731–742.
- Reina, Andreagiovanni, Thomas Bose, Vito Trianni, and James A. R. Marshall. 2018. “Psychophysical Laws and the Superorganism.” *Scientific Reports* 8 (1): 1–8.
- Rendell, Luke, Robert Boyd, Daniel Cownden, Marquist Enquist, Kimmo Eriksson, Marc W. Feldman, Laurel Fogarty, Stefano Ghirlanda, Timothy Lillicrap, and Kevin N. Laland. 2010. “Why Copy Others? Insights from the Social Learning Strategies Tournament.” *Science* 328 (5975): 208–213.
- Reynolds, Andy M., Patrick Schultheiss, and Ken Siu-Kei Cheng. 2013. “Are Lévy Flight Patterns Derived from the Weber-Fechner Law in Distance Estimation?” *Behavioral Ecology and Sociobiology* 67 (8): 1219–1226.
- Reynolds, Craig W. 1987. “Flocks, Herds and Schools: A Distributed Behavioral Model.” *ACM SIGGRAPH Computer Graphics* 21 (4): 25–34.
- Rubenstein, Michael, Christian Ahler, and Radhika Nagpal. 2012. “Kilobot: A Low Cost Scalable Robot System for Collective Behaviors.” In *2012 IEEE International Conference on Robotics and Automation*, 3293–3298. New York: IEEE.
- Russell, Stuart J., and Peter Norvig. 2016. *Artificial Intelligence: A Modern Approach*. London: Pearson.
- Saffre, Fabrice, Hanno Hildmann, and Jean-Louis Deneubourg. 2018. “Can Individual Heterogeneity Influence Self-Organised Patterns in the Termite Nest Construction Model?” *Swarm Intelligence* 12 (2): 101–110.
- Santana, Pedro, and Luís Correia. 2010. “A Swarm Cognition Realization of Attention, Action Selection, and Spatial Memory.” *Adaptive Behavior* 18 (5): 428–447.
- Sasaki, Takao, and Stephen C. Pratt. 2018. “The Psychology of Superorganisms: Collective Decision Making by Insect Societies.” *Annual Review of Entomology* 63:259–275.
- Sayama, Hiroki. 2013. “PyCX: A Python-Based Simulation Code Repository for Complex Systems Education.” *Complex Adaptive Systems Modeling* 1 (1): 1–10.
- Schmickl, Thomas, and Heiko Hamann. 2011. “Beeclust: A Swarm Algorithm Derived from Honeybees.” In *Bio-inspired Computing and Communication Networks*, edited by Yang Xiao, 95–137. Boca Raton, FL: CRC Press.
- Schmickl, Thomas, and Istvan Karsai. 2016. “How Regulation Based on a Common Stomach Leads to Economic Optimization of Honeybee Foraging.” *Journal of Theoretical Biology* 389:274–286.
- Schmickl, Thomas, Christoph Möslinger, and Karl Crailsheim. 2006. “Collective Perception in a Robot Swarm.” In *Swarm Robotics*, edited by Erol Şahin, William M. Spears, and Alan F. T. Winfield, 144–157. Berlin: Springer.
- Sih, Andrew, and Marco Del Giudice. 2012. “Linking Behavioral Syndromes and Cognition: A Behavioral Ecology Perspective.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1603): 2762–2772.
- Silva, Fernando, Paulo Urbano, Luís Correia, and Anders Lyhne Christensen. 2015. “odNEAT: An Algorithm for Decentralised Online Evolution of Robotic Controllers.” *Evolutionary Computation* 23 (3): 421–449.
- Thenius, Ronald, Daniel Moser, Joshua Cherian Varughese, Serge Kernbach, Igor Kuksin, Olga Kernbach, Elena Kuksina, et al. 2016. “subCULTron—Cultural Development as a Tool in Underwater Robotics.” In *Artificial Life and Intelligent Agents Symposium*, 27–41. Cham, Switzerland: Springer.
- Thierry, Bernard, Guy Theraulaz, Jean-Yves Gautier, and B. Stiegler. 1995. “Joint Memory.” *Behavioral Processes* 35 (1–3): 127–140.
- Trianni, Vito, Elio Tuci, Kevin M. Passino, and James A. R. Marshall. 2011. “Swarm Cognition: An Interdisciplinary Approach to the Study of Self-Organising Biological Collectives.” *Swarm Intelligence* 5 (1): 3–18.
- Valentini, Gabriele, Heiko Hamann, and Marco Dorigo. 2014. “Self-Organized Collective Decision Making: The Weighted Voter Model.” In *Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '14)*, 45–52. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Valentini, Gabriele, Heiko Hamann, and Marco Dorigo. 2015. “Efficient Decision-Making in a Self-Organizing Robot Swarm: On the Speed versus Accuracy Trade-Off.” In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15)*, 1305–1314. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Van den Driessche, Charlotte, Françoise Chevrier, Axel Cleeremans, and Jérôme Sackur. 2019. “Lower Attentional Skills Predict Increased Exploratory Foraging Patterns.” *Scientific Reports* 9 (1): 1–7.
- Wahby, Mostafa, Mary Katherine Heinrich, Daniel Nicolas Hofstadler, Ewald Neufeld, Igor Kuksin, Payam Zahadat, Thomas Schmickl, Phil Ayres, and Heiko Hamann. 2018. “Autonomously Shaping Natural Climbing Plants: A Bio-hybrid Approach.” *Royal Society Open Science* 5 (10): 180296.

- Wahby, Mostafa, Julian Petzold, Catriona Eschke, Thomas Schmickl, and Heiko Hamann. 2019. "Collective Change Detection: Adaptivity to Dynamic Swarm Densities and Light Conditions in Robot Swarms." In *Proceedings of the ALIFE 2019: The 2019 Conference on Artificial Life*, edited by Harold Fellersmann, Jaime Bacardit, Ángel Goñi-Moreno, and Rudolf M. Füchslin, 642–649. Cambridge, MA: MIT Press.
- Werfel, Justin, Kirstin Petersen, and Radhika Nagpal. 2014. "Designing Collective Behavior in a Termite-Inspired Robot Construction Team." *Science* 343 (6172): 754–758.
- Wilson, David Sloan, and Elliott Sober. 1989. "Reviving the Superorganism." *Journal of Theoretical Biology* 136 (3): 337–356.
- Winfield, Alan F. T., and Julien Nembrini. 2006. "Safety in Numbers: Fault Tolerance in Robot Swarms." *International Journal on Modelling Identification and Control* 1 (1): 30–37.
- Wolf, Max, Ralf H. J. M. Kurvers, Ashley J. W. Ward, Stefan Krause, and Jens Krause. 2013. "Accurate Decisions in an Uncertain World: Collective Cognition Increases True Positives While Decreasing False Positives." *Proceedings of the Royal Society B: Biological Sciences* 280 (1756): 20122777.
- Yasuda, Toshiyuki, Akitoshi Adachi, and Kazuhiro Ohkura. 2014. "Self-Organized Flocking of a Mobile Robot Swarm by Topological Distance-Based Interactions." In *2014 IEEE/SICE International Symposium on System Integration*, 106–111. New York: IEEE.

6 Soft Robotics: A Developmental Approach

Luca Scimeca and Fumiya Iida

6.1 Introduction

In this chapter we will first introduce and review soft robotics research, with emphasis on how compliance and softness have changed the robotics landscape in the past two decades. We will then briefly discuss the key ideas in developmental robotics that are fundamental for understanding the relationship between biological and artificial systems, and examine how the developmental sciences and soft robotics are irrevocably linked, into what we have chosen to name “developmental soft robotics.” Here, in fact, the two fields can be merged into one in which the developmental sciences can aid in the design and make of soft robots that can then be used as platforms to better understand biological systems. We will finally discuss how phylogenetic development, ontogenetic development, and short-term adaptation are indeed naturally suited to be embedded within a “soft” robotic context. (For further reading, see Trivedi et al. 2008; Pfeifer, Iida, and Lungarella 2014; Laschi et al. 2016.)

6.2 Bioinspired Soft Robotics

Deformation is a fundamental characteristic of biological systems. Almost 90 percent of the human body is composed of soft tissue; many vital organs such as the heart, lungs, muscles, eye lenses, and more depend on deformation of materials.

In bipedal walking, for example, evidence has shown how the soft tissue of the body might not only cushion the impacts of each stride, but also save muscles the effort of actively dissipating energy, while performing a considerable amount of the total positive work, per stride, by soft tissue elastic rebound (Zelik and Kuo 2010).

In the past few decades, there has been an unprecedented advancement in material science and manufacturing techniques, furthering our knowledge of functional materials and empowering artificial systems with newfound capabilities. These advancements, together with a better understanding of biological systems, have given rise to the era of soft robotics, in which bioinspired robotics platforms make use of soft and deformable materials to achieve more flexible, adaptable, and robust behaviors (Kim, Laschi, and Trimmer 2013; Hughes et al. 2016).

Since the dawn of soft robotics, the application of material science and soft-body compliance has changed the robotics landscape. In manipulation, for example, the “universal gripper,” a soft gripper capable of particle jamming through vacuum pressure control, has been shown to be able to grasp a large number of objects (Brown et al. 2010). Other solutions for grasping and manipulation range from tentacle-like systems (Laschi et al. 2012) to pneumatic soft grippers (Yap, Ng, and Yeow 2016) and human-inspired soft-robotic hands (Hughes, Maiolino, and Iida 2018; figure 6.1).

Animal-inspired soft robots are among the most developed subareas of soft robotics, where the robot platforms range from worms (Seok et al. 2010) or caterpillars (Lin, Leisk, and Trimmer 2011) to octopuses (Laschi et al. 2012), fish (Katzschmann et al. 2018), and others (figure 6.1). In wormlike soft robots, for example, akin to their biological counterparts, the contraction of longitudinal muscles followed by the contraction of circumferential muscles simulates a traveling wave along the body, generating locomotion (Trueman 1975). In caterpillars, motion is generated by coordinated control of the time and location of the prolegs attachment to the substrate, together with waves of muscular contraction (Belanger and Trimmer 2000).

The ability to mimic these unique systems makes soft robots an exciting new field, where the limits of the (rigid) robots of the past century can be overcome with newfound solutions.

6.2.1 Soft Materials and Soft Actuation

The area of soft robotics is inevitably connected to the field of material science, in which new discoveries in the latter facilitate progress in the former. For a soft robot to be able to use material compliance to aid in robotics tasks, it is necessary for the make of the robot to be, at least in part, deformable. Elastomeric (polymer) materials, like EcoFlex or Drag-




Figure 6.1

Bioinspired soft robot examples. (a) Worm-inspired soft robot. *Source:* Seok et al. 2010. (b) Caterpillar-inspired soft robot. *Source:* Lin et al. 2011. (c) Octopus-inspired tentacle. *Source:* Cianchetti et al. 2011. (d) Human-inspired soft passive hand. *Source:* Hughes et al. 2018. (e) Fish-inspired soft robot. *Source:* Katzschmann et al. 2018.

onSkin (Siegenthaler et al. 2011), have been at the center of researchers' attention for several years, with new substances being discovered every year. Moreover, the advent of three-dimensional printing technology has led to much faster robot design and testing operations than before, facilitating rapid and cheap prototyping in soft robotics.

Actuation poses one of the biggest challenges. In many animals, the coaction of a large number of muscles distributed over the body is capable of generating relatively high forces, facilitating coordinated and robust action. Replicating this ability is no easy feat, as the majority of the robotics solutions lack the ability to generate forces comparable to the industrial robots of the past.

Four main soft-actuation techniques currently exist: tendon driven, pressurized air or fluids, dielectric elastomeric actuators, or DEAs, and shape memory alloys, or SMAs (Kim et al. 2013). First, tendon-driven actuation mimics biological musculoskeletal systems, in which actuation is achieved through the pull and release of tendons, via the appropriate control of motors (figure 6.2*a*). Although a powerful and widespread actuation technique, a large number of tendons are usually necessary to achieve complex behaviors, and control complexity increases along with the number of motors necessary to control the tendons. For softer robots, like continuum soft robots, this type of actuation usually does not scale. Second, the employment of fluids is one of the most powerful actuation techniques for soft robots, capable of generating high forces and displacements. The actuation usually consists of varying the pressure inside predesigned chambers within the body of the robot to achieve their expansion and contraction and generate motion or morphological changes (figure 6.2*b*). However, these actuation systems are usually bulky and heavy and require high power sources, making them unsuitable for untethered robotics systems (Laschi and Cianchetti 2014). Third, DEAs are made of soft materials that can be actuated through electrostatic forces (figure 6.2*c*). DEAs have been shown to have high-strain/stress and mass-specific power. However, the need for DEAs to be prestrained imposes rigid constraints on the robots' design (O'Halloran, O'Malley, and McHugh 2008). Finally, SMAs, with the most common nickel titanium alloys, can generate force through a change in shape due to a rise or fall in the temperature of the material (figure 6.2*d*). Temperature change control, however, is a challenge. High voltages are usually required to achieve temperature changes, and robustness over varying temperatures in the environment is still an issue to be overcome (Rodrigue et al. 2017). Other methods exist; it is possible, for example, to induce pneumatic contraction by evaporating ethanol via resistive heating (Miriyeve, Stack, and Lipson 2017) or to achieve bending through combustion (Tolley et al. 2014). Other issues, such as reduced output force or slow speed, however, come into play (Rich, Wood, and Majidi 2018). Soft robotics actuation and material sciences are still an ever-changing field, with new solutions being expedited by fast prototyping and iteration.

6.2.2 Soft Robot Control, Simulation, and Learning

Soft-robotic control poses several challenges and opportunities. Here, the "degree of softness" matters. Take, for example, a rigid robotic hand with the palms and fingertips covered with an elastomeric material. The control of the hand is usually possible to achieve with classical methods (i.e., inverse kinematics), in which the complexity of the control depends on the complexity of the mechanical system. If the hand were entirely rigid, achieving the


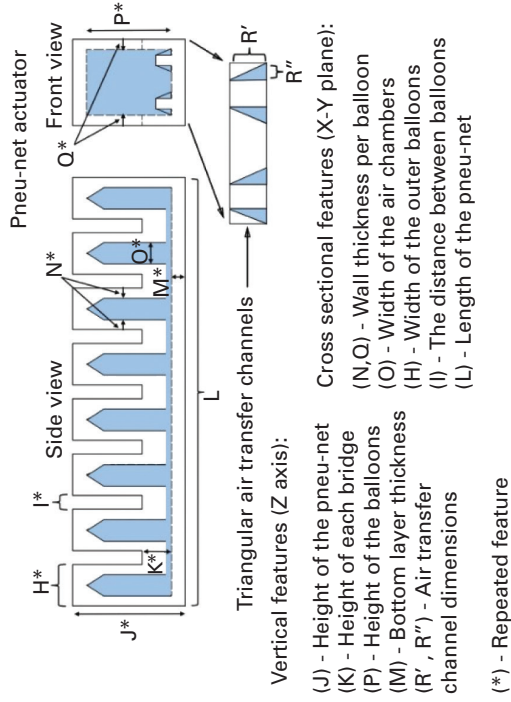







Figure 6.2

Examples of some of the main actuation mechanisms used for soft robotic systems. (a) Tendon-driven continuum robot and model. *Source:* Rucker and Webster 2014; Geng et al. 2018. (b) Pneumatic soft actuator. *Source:* Yirmibesoglu et al. 2018. (c) Variable stiffness dielectric elastomer actuator. *Source:* Shimtaka et al. 2015. (d) Curved memory alloy-based soft actuator. *Source:* Rodrigue et al. 2017.



d



c




Figure 6.2 (continued)

appropriate control to perform a “light” touch might not be trivial. By appropriately exploiting the mechanical passive dynamics of the soft fingers, the complexity of the control can be reduced to achieve the desired grasping behavior, averting the need for submillimeter precision in robot control (Pfeifer, Lungarella, and Iida 2007; Iida and Laschi 2011). However, as the “degree of softness” in the body increases, new challenges arise.

A robot made entirely of elastomeric materials—for example, one emulating the tentacle of an octopus or the trunk of an elephant—cannot be controlled classically; moreover, proprioception and simulation become problematic. As opposed to the hard links with sliding or rotational joints in classical robots, the continuity and softness of the body makes the control and simulation of continuous soft robots much more difficult. Novel actuation methods aid robotics researchers in their endeavors to achieve desired robot control (section 6.2.1), and new sensing and control methods are discovered on a daily basis (Rus and Tolley 2015). To achieve autonomy and go beyond open-loop control for soft robots, both proprioception and tactile sensing are required.

Much effort has been put into the sensorization of soft robots. The most common soft sensors are perhaps strain sensors, which are soft, deformable sensors capable of sensing body deformations through stretching. It is thus possible to embed such sensors into the (soft) body of a robot without influencing its ability to deform. Some of the most widespread sensors are based on resistive (Homberg et al. 2015) or capacitive (Maiolino et al. 2015) technologies. Recently, work in Galloway et al. (2019) and Scimeca et al. (2019) have shown how it is possible to achieve a high-fidelity proprioceptive understanding of a continuum soft body through sensorization via fiber-optic and capacitive tactile sensors, respectively.

In the context of control and simulation, learning plays a fundamental role. With the infinite degrees of freedom posed by a continuum soft body, for example, precise control via classical methods is hard and usually does not scale. Model-based solutions relying on the piecewise constant curvature assumption have been shown to work for small, tentacle-like robots (Della Santina et al. 2018). However, the error in the controller always increases with an increase in the number of soft segments within the robots. The models, in fact, are usually too simplistic to accurately capture the complexity of continuum soft robots. Learning in this case has been shown to be useful in compensating for a lack of knowledge or model complexity (Scimeca, Maiolino, and Iida 2018, 2020; Rosendo, von Atzigen, and Iida 2017).

6.3 Developmental Soft Robotics

Cognitive developmental robotics (CDR) is an area of research in which robotics and the developmental sciences merge into a unique field, one that seeks to better robotics with insights from developmental sciences and further our understanding of developmental sciences through the use of robotics platforms (Lungarella et al. 2003). The need for CDR to be a research area on its own arose at the dawn of the twenty-first century from the need to understand not only the cognitive and social development of individuals, as explored in the area of epigenetic robotics (Zlatev and Balkenius 2001), but also the acquisition and development of motor skills and how they, as well as morphology, influence the development of higher-order cognitive functions (Lungarella et al. 2003; Asada et al. 2001, 2009).

In this context, robots can be used as experimental subjects, where developmental models can be implemented in robotics platforms, and scientists can gain insights from behavioral analysis, an approach known as synthetic methodology (Scheier and Pfeifer 1999; Sporns 2003).

In stark contrast to the traditional computationalist approach, in developmental robotics there is no clear separation between the physical body, the processes that determine reasoning and decision-making (cognitive structure), and the symbolic representation of entities in the world. Rather, these processes influence each other, and intelligence emerges from their interaction. Developmental robotics is treated in detail in chapter 3.

One of the most difficult tasks in modern-day robotics is to achieve an appropriate robot design for a robot to perform certain tasks in the world. The advent of soft robotics, if anything, has increased the complexity of robots, revoking the rigidity constraints established in the earlier century and bringing about a new era. In this new era, robot design is driven by factors much like biological systems, in which functional morphology, coordinate sensorimotor action, physical adaptation, and embodiment all contribute to the “robot’s survival” in the world and to its ability to see a task to completion.

Developmental soft robotics aims to bring together the areas of soft robotics with those of developmental robotics and the developmental sciences. These, in fact, are irrevocably linked, as we will show.

6.3.1 Soft Robotics and Developmental Timescales

Within the developmental sciences, in its simplest form, the development of a biological organism can be distinguished on three different scales: phylogenetic, ontogenetic, and short-term.

In biological organisms, *phylogenetic development* has the largest timescale, in which changes happen at the level of groups of organisms, over many generations, and processes such as natural selection are responsible for certain “traits” surviving and evolving, while others become extinct. Akin to phylogenetic development is soft robotics design, in which the design of robots is adaptive and ever changing to comply and conform to the task the robot must achieve. Currently, much of the adaptation is due to human design and biased by human skill and experience. However, new methodologies for autonomous designs are a hot research topic, and processes such as evolutionary algorithms have shown promise in the past (Nolfi and Floreano 2000; Doncieux et al. 2015).

Ontogenetic development concerns changes throughout and within the life span of an organism and includes growth and bodily adaptation. The ability of robots to “morph” throughout their life span to achieve desired behaviors has been one of the key advantages of soft robots, as opposed to their rigid counterparts of the previous century. Robots navigating through growth like fungal hyphae (Hawkes et al. 2017), elongating their bodies due to pressure and changing their stiffness to alter their body dynamics and achieve different behaviors (Cianchetti et al. 2013), are examples of such adaptability.

Short-term adaptation refers to the shortest adaptive and developmental timescale of all, in which adaptation needs to be achieved instantaneously. Short-term adaptation is perhaps the most naturally suited to be discussed in a soft setting. In the past this type of adaptation needed to be actively achieved at the control level, where real-time control

would allow short-term adaptive behavior through mechanical or sensory feedback. Within the soft robotics framework, much like biological organisms, the short time adaptation is just a consequence of the soft, instantaneous deformation of the soft body itself. When we delicately slide our finger through a ridged surface, for example, the need for complex and precise control is voided by the ability of our dermis to deform and conform to the surface under our touch. Much like the illustrated example, the compliance and softness of materials, in soft robots, can achieve short-term adaptation. The mechanical feedback becomes only a physical consequence of contact, and compliance can naturally suppress the need for complex controllers. Figure 6.3 illustrates the main idea behind the developmental soft robotics framework.

6.3.2 Functional Morphology and Morphological Computation

When designing robotics systems, if shape was initially the most salient of morphological features, with the advent of soft robotics this may no longer be the case. Materials at different levels of elasticity have demonstrated the ability to perform “computation” (Scimeca et al. 2018; Eder, Hisch, and Hauser 2018). Recent work in Scimeca et al. (2018), for example, has shown how complex haptic information can be used to classify objects based on different properties, solely based on clustering analysis. The simplicity of the inference is possible due to a “soft filter” or elastic layer between the tactile sensor and the object. When changing the properties of the elastic layer, the tactile information is appropriately influenced (spatially filtered) in order to induce object similarities with respect to different object properties, like edges or elongation. The “intelligence” is here in the body, since

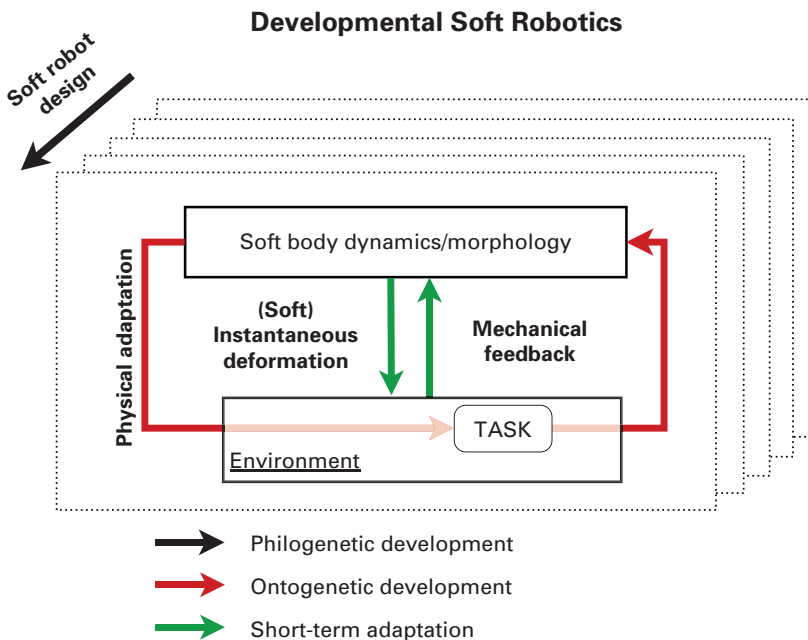


Figure 6.3
Developmental soft robotics.

the body's ability to appropriately mold the sensory information allows for the agent's higher cognitive functions to solve the object classification problem with simple clustering methods, without prior training or supervision, an otherwise impossible feat.

A paradigm trying to make use of the complex body-environment interactions is the "reservoir-computing" framework of computation. The original idea behind reservoir computing begins with network computation, in which an input is fed to a network, which computes a corresponding output. In reservoir computing, a fixed random dynamical system, also known as a reservoir, is used to map input signals to a higher-dimensional space. The "readout" final part of the network, then, is trained to map the signals from the higher-dimensional space to their desired output. As previously mentioned, soft robots, as well as biological organisms, are usually made, at least in part, of soft materials. The body dynamics of soft robots are thus very complex, highly nonlinear, and high dimensional, making control challenging. Through the reservoir-computing paradigm, it is possible to capitalize on the complexity of such a system by exploiting the soft body as a computational resource, using the body dynamics to emulate nonlinear dynamical systems, and, as a result, off-loading some of the control to the body itself (Nakajima et al. 2013, 2015). Nakajima et al. (2014), for example, have shown it is possible to control a complex continuum soft arm, inspired by the tentacle of an octopus, in a closed loop without any external controller, by using the body of the robot as a computational resource. In this light, high nonlinearity and complexity may be a desirable property of the body, and design might have to be thought of accordingly.

An additional property that allows soft bodies to be used as a computational resource is memory. The soft body dynamics of soft robots, in fact, can exhibit short-term memory, allowing robots to emulate functions that require embedded memory (Nakajima et al. 2014). When underactuating a continuum soft robot, for example, it may be that the control mechanism is not deterministic with respect to the behavior of the robot. In these cases the behavior of the robot may depend not only on the induced control and its current state but also on the history of the previous robot states, as it may be the case when actuating a soft tentacle arm by moving one of its extremities.

6.3.3 Emergent Behaviors of Soft Robots

At the dawn of the twenty-first century, the concept of "morphofunctional machines" was proposed. Morphofunctional machines were defined as those that were adaptive by being able to change their morphology as they performed tasks in the real world (Hara and Pfeifer 2003). In this context, changes at different timescales were already argued to be important—that is, short-term, ontogenetic, and phylogenetic, or evolutionary. It is important to note that the adaptation and the resolution of the task here is achieved not at the control level but at the morphological level.

As advocated by the developmental robotics paradigm (chapter 3), intelligence and coordinated action are the result of complex interactions between the body, the mind, and the environment. The latter, in fact, plays an important role in determining the behaviors of the artificial or natural organisms living within it.

One of the most influential experiments of the last two decades was the "dead fish experiment," performed in collaboration with Harvard and the Massachusetts Institute of

Technology (MIT) in 2005 (Beal et al. 2006). In the experiment, a dead fish was able to swim upstream even when its brain was clearly sending no control impulse. Upon further study it was apparent how the streamlined body of the fish, passively oscillating, was capable of turning the surrounding energy into mechanical energy and thus propel itself forward passively. Although the morphology and make of the body allowed the dead fish to transduce the surrounding energy, the environment was the enabling factor. The vortices created by water streams were key in the experiment, as they generated the energy to be transduced and recreated the conditions for the body to manifest its propelling abilities. The interaction between the body and the environment were, in fact, the decisive factors in determining the observed behavior. A similar influential experiment was the passive dynamic walker. The make of the robot, with kneecaps, springs, pendulum-like leg swings, and more, was capable of stable, humanlike, and low-energy bipedal locomotion without any complex control. However, the environment initiated and stabilized the walking locomotion, as it manifested when the robot was placed on a downward slope (Collins et al. 2005), allowing the potential energy to be skillfully turned into kinetic energy.

In robot design it is therefore always necessary to take the environment into account. Much like the examples previously mentioned, the body and the brain are often not enough to achieve useful objectives. Things in the world exist to affect and change their surroundings and live within the environment they are situated in (Mataric 2006). In this context it is in the interplay of the body and the environment that intelligent, situated behavior can be observed and that morphology can be empowered and purposefully adapted.

6.3.4 Sensing and Perception of Soft Robots

In nature, morphology plays a fundamental role within the sensing landscape, mechanically converting, filtering, and amplifying sensor stimuli from the outside world to make sense of the surrounding environment or internal states (Towal et al. 2011; Iida and Nurzaman 2016). In rats and mice, for example, vibrissae, or sensitive tactile hairs, have been known to confer to these mammals specialized tactile capabilities, aiding them in a number of sensory discrimination tasks (Prescott et al. 2009). In a similar manner, most mammals have evolved to mediate vision through compound eyes, compromising resolution for larger fields of view and high temporal resolution, and enabling fast panoramic perception (Land and Nilsson 2012). Within the biomimetic robotics field, attempts have been made to endow robotics systems with the capabilities of organisms observed in nature. Haptic robot perception through whiskers (Pearson et al. 2011) and compound vision (Floreano et al. 2013) are two such examples (figure 6.4).

Soft sensing is one of the most popular fields within the soft robotics landscape. Augmenting soft robotics systems with the ability to sense the environment can enable robots to react to unknown events, to improve their control and morphology over time, and to capture information or reason about entities in the world. Sensorizing soft robots is no easy task. One of the goals within this field is to devise sensors that themselves exhibit some “soft” behavioral characteristics; usually, flexibility (i.e., can be bent) and stretchability (Lu and Kim 2014) are desirable. Currently, approaches to achieve stretchable electronics include wavy circuits (Majidi 2014; Rogers, Someya, and Huang 2010) and conductive liquids (Cheng and Wu 2012). One of the most widespread soft sensors are strain sensors, shown to be highly

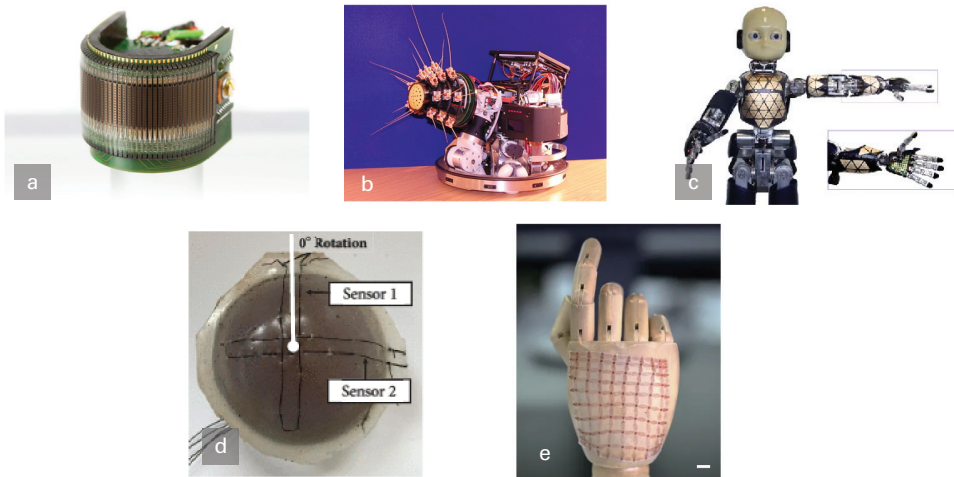


Figure 6.4

Bioinspired flexible and soft sensing examples. (a) Artificial compound eyes. *Source:* Floreano et al. 2013. (b) Robotic tactile vibrissal sensing. *Source:* Pearson et al. 2011. (c) iCub robot with large-area flexible capacitive tactile skin. *Source:* Hoffmann et al. 2017. (d) Conductive thermoplastic elastomer sensorized universal gripper. *Source:* Hughes and Iida 2017. (e) Stretchable and conformable sensor for multinational sensing. *Source:* Hua et al. 2018.

elastic (Muth et al. 2014). New embedding methodologies have also demonstrated the possibility of embedding strain sensors within elastomers through three-dimensional printing techniques. Other flexible sensing technologies such as capacitive tactile sensing (Maiolino et al. 2013) and fiber optics (Galloway et al. 2019) have been used within soft robotics systems.

As previously mentioned, sensorimotor coordination and morphology can enhance the sensing capabilities of robotics systems. Sensors should not be thought of simply as independent and self-sufficient technologies. Instead, it is fundamental to think of sensor technologies as apparatuses that reside within a body. The body dynamics derived from its morphological properties, coupled with the environment the robotic system is situated in, should all contribute to the sensor morphology, its characteristics, and its perceptual capabilities. The appropriate coupling of these factors has been shown to improve the sensing capabilities of robotic systems (Iida and Pfeifer 2006). In Hughes and Iida (2017), for example, the sensorization of a universal gripper was achieved with a pair of conductive thermoplastic elastomer (CTPE) strain sensors (figure 6.4d). Differential sensing was then used to compute deformations within the soft body. Morphology, however, was key. By weaving the strain sensor in different patterns within the soft gripper, information regarding the magnitude, orientation, or location of a deformation could be detected. Because such sensing is also inescapably linked to motor control, mechanical dynamics, and the objectives of the robotic system, the concept of “adaptive morphology” has recently been proposed (Iida and Nurzaman 2016), wherein the iterative design, assembly, and evaluation of sensor methodologies attempt to explain the adaptive nature of the perceptual abilities of living organisms.

6.3.5 Adaptation and Growth

The principles previously discussed encourage a different approach to design, in line with endowing robots with the ability to adapt to ever-changing environments and indeed make use of the environment as a means of solving their assigned tasks. Besides design principles at a phylogenetic scale and instantaneous deformation on the short-term scale via material properties and design, another important factor is ontogenetic change and adaptation. Plants, for example, are capable of continuously changing their morphology and physiology in response to variability within their environment in order to survive (Mazzolai, Beccai, and Mattoli 2014). Inspired by the unique abilities of plants to survive in diverse and extreme environments, a stream of researchers have more avidly tried to reproduce some of their adaptivity in robotics systems. Plantoids, or robotic systems equipped with the distributed sensing, actuation, and intelligence to perform soil exploration and monitoring tasks, have started to gain traction in this direction (Mazzolai, Beccai, and Mattoli 2014). Rootlike artificial systems in Sadeghi et al. (2013) and (2014), for example, have been shown to be able to perform soil exploration through novel methodologies simulating growth via elongation of the robot's tip. Other plant-inspired technologies in biomimicry and the material sciences include Velcro, from the mechanisms behind the hooks of plant burrs (Velcro SA 1955), bamboo-inspired fibers for structural engineering materials (Li et al. 1995), and novel actuation mechanisms in Taccola et al. (2013) based on reversible adsorption and desorption of environmental humidity and, in Mazzolai et al. (2010), based on the osmotic principle in plants.

Another important factor in ontogenetic adaptivity is the ability of organisms to mend their own tissue over their life spans. Endowing artificial systems with self-healing abilities has recently become of primary importance, setting the landscape for untethered robots to “survive” for longer periods of time in more uncertain and dynamic task environments. Self-healing of soft materials is typically achieved through heat treatment of the damaged areas, which allow some polymers to reconnect and retrieve most of their structural properties. In (Terryn et al. 2017), for example, a soft gripper, a soft hand, and artificial muscles were developed with Diels-Alder materials (Scheltjens et al. 2013). In the developed systems, the Diels-Alder materials were shown to be reversible at temperatures of 80°C, recovering up to 98 to 99 percent of the mechanical properties of the polymers postdamage.

6.3.6 Tool Use and Extended Phenotype

In biology, the phenotype is the set of observable traits of an organism, including its morphology, developmental process, and physiological properties. The idea of extended phenotypes was first introduced by Richard Dawkins (1982) when he argued that the original concept of phenotype might have been too restricted. In fact, the effects that a gene may have are not limited to the organism itself but to the environment the organism is situated in, through that organism's behavior. The coupling of an artificial agent and its environment was discussed in section 6.3.3. The extended phenotype notion, however, extends to even more radical concepts.

One of the most fascinating examples of this is found in primates, corvids, and some fish, which have been found to purposefully make and use “tools” to achieve goals within

their environments, such as the acquisition of food and water, defense, recreation, or construction (Shumaker, Walkup, and Beck 2011).

Extending the phenotype concept, the observable traits of the organisms should be augmented to include their extended functionalities, behaviors, and morphology, as derived from the use of the tool in question. When a primate is holding a small branch, for example, the physical characteristics of the primate are undeniably changed: its reach is longer, and its weight and morphology are affected, as is its stance to balance on two or three limbs or its ability to affect the environment around it. Under the extended phenotype concept, these changes must be captured within the phenotypic traits of the organism.

In the context of soft developmental robotics, the ontogenetic development of robotics systems should include their ability to adapt to their environments over their life span (physical adaptation) and indeed their ability to augment their functionality by the active creation and use of tools initially excluded from their phenotypic traits. This ability was previously investigated in Hoffmann et al. (2010) and Nabeshima, Kuniyoshi, and Lungarella (2006), where it was obvious that at the foundation of the idea of tool use was the concept of body schema (cf. chapter 3). The body schema in this scenario requires adaptability and alterability throughout ontogenetic development to cope with the changes in one's body, including growth, as well as with the extended capabilities conferred by the use of tools. An understanding of the tool is necessary here (Wang, Brodbeck, and Iida 2014). Nabeshima, Kuniyoshi, and Lungarella (2006) argued that the temporal integration of multisensory information is a plausible candidate mechanism to explain tool use incorporation within the body schema. Another core component in this context is proprioceptive sensing, or the ability to sense self-movement and body position. Proprioception also plays a significant role in the perception/action model of body representations (de Vignemont 2010).

6.4 Conclusion

Throughout this chapter we have examined the various aspects of bioinspired robotics, with emphasis on soft robotics and the idea that intelligence is exhibited as an interplay, and reciprocal dynamical coupling, of the brain, the body, and the environment. The concept of developmental soft robotics was introduced in this context, in which some design principles can be established on three different timescales, aiding and enabling roboticists and researchers to develop systems for a new generation of robots. Many enabling technologies for sensing and actuation have driven progress in the past few decades and have allowed robots to pass from rigid and industrial to soft and human-friendly. These robots have been shown to achieve locomotion, to pick up and manipulate objects, to safely interact with humans, and much more. However, many challenges still await this field, as the road to the ultimate goal of creating machines with abilities akin to those of organisms in the animal world is only in its early stages.

6.4.1 Physical Soft Robot Evolution

On the phylogenetic timescale, the question of how to achieve complex embodied behavior has been answered by nature for a very long time. The concept of evolution in biological

organisms is fairly straightforward, where evolution is thought of as the change in inheritable characteristics of populations over successive generations (Hall and Strickberger, 2008). Due to various sources of genetic variation, new generations have increasingly different traits, and via a mediating process like that of natural selection, some traits will ensure higher or lower chances of survival (Scott-Phillips et al. 2014). Eventually, the surviving population has all the different traits that we can now see in the immense variety of living organisms on our planet, which have adapted to use a plethora of different methodologies and techniques to ensure their survival.

The field of phylogenetics in the context of soft robotics is tightly coupled with this concept, and consequently, this field has a major impact on emergent design and control in robotics. In the area of “evolutionary robotics,” evolutionary computation is used to develop physical designs or controllers for robots (cf. chapter 4). Evolutionary computation takes inspiration from biological evolution. In robotics, for example, it is possible to create an initial set of candidate robots and encode their physical and or control characteristics numerically. By testing the robot population against a specific task, it is then possible to identify which combination of morphology and control performed better. The encoded characteristics of the best-performing robots can then be perturbed and used to create a new generation of robots that can be tested again. The iteration of this process for thousands of iterations has been shown to achieve robust controls (Mautner and Belew 2000; Fleming and Purshouse 2002) and designs (Lund, Hallam, and Lee 1997; Lipson and Pollack 2000; Pfeifer, Iida, and Bongard 2005; Vujovic et al. 2017; Brodbeck, Hauser, and Iida 2015).

One of the biggest limitations of evolutionary algorithms lies with the resources and time necessary to achieve good controllers or designs. Because the iteration of robot design, robot testing, and robot evaluation are very time-consuming, it is generally not feasible to apply evolutionary algorithms in very complex problems by starting from a generic, nonbounded, encoding of robot characteristics. The world of simulation has historically been more suited for evolutionary algorithms (Lipson and Pollack 2000; Mautner and Belew 2000; Nolfi et al. 1994) given the ease with which populations can be created, tested, and iterated over. The controllers and designs found, however, are usually not robust real-world solutions, as simulation environments are still very limited, and the solutions found within them do not necessarily correspond to solutions in the real world (Jakobi, Husbands, and Harvey 1995). Moreover, depending on the complexity of the problem, computational resources are still an issue.

In soft robotics, given the complexity of the bodies and the interactions emerging from them, design and control pose two of the biggest problems. Evolutionary algorithms find themselves suited as a candidate solution, but the limitations previously mentioned still apply. Further advancements in virtual reality engines, new manufacturing methods for fast prototyping, advancements in material science, and the ever-increasing power of computing, however, may bypass some of these limitations in the near future.

6.4.2 Complexity and Scalability

As of today, the robots we see still “feel” unnatural; they move slowly and sluggishly; humanoid robots still do not possess the ability to walk, run, or move the way humans do; they cannot reason about the world the same way we do and they get confused when unknown events occur (Pfeifer, Lungarella, and Iida 2012). One of several reasons con-

tributing to this fact is complexity. The number of actuators and distributed sensors present in humans is much too high to be replicated by motors and standard sensors in machines. This complexity poses a problem, as does controlling the coupling of a high number of motors and sensors. Even when dealing with subproblems, like humanoid hands, the complexity may very well already be too high to try and tackle with standard methods. Some attempts to replicate complexity were made, for example, by replicating in a robotic manipulator many of the degrees of freedom present in a human hand (Tuffield and Elias 2003). This approach, however, did not give the results many were hoping for, as complexity in the body was coupled with complexity in the control, and achieving an adaptable, smooth grasp and manipulation behavior was no easy task. Recent advances have shown how an underactuated, or even passive, hand can achieve complex behaviors, if its interactions with the environment are appropriately exploited (Hughes et al. 2016, 2018). It is here that complexity can be displaced, since complex behavior can emerge from simple design when appropriate interactions take place.

Within this framework, many questions remain. It is, in fact, unclear how design should be achieved to avoid or exploit complexity. Exploiting environmental constraints is no easy feat, as the constraints to be exploited are also tightly coupled with the task at hand. In soft robotics the make of the robots themselves leads to highly nonlinear behaviors and robots with complex dynamics. Paradigms like that of reservoir computing can capitalize on the complexity of such structures, using them as a computational resource and thus making complexity a desirable feature. Control, however, is still hard to achieve, and mathematical models fail to comprehensively account for dynamical interactions when the complexity of the body becomes too high. This complexity presents infinite challenges and opportunities, which the ever-changing landscape of robotics will have to face in the near future.

6.4.3 Learning through the Body

The advancements in artificial intelligence (AI) in the last two decades have begun a scientific revolution, endowing machines with the possibility to achieve superhuman performance levels in several different fields, like image-based object detection (Schmidhuber 2015), virtual agent control (Mnih et al. 2015), and haptic texture identification (Fishel and Loeb 2012). In robotics, machine learning has been extensively used both on the perceptual side, such as for object detection and recognition, and on the control side, such as for robot trajectory planning and motor control.

The most powerful machine-learning algorithms make use of supervision, or the knowledge of target labels, to improve performance over time or trials. Broadly speaking, from the machine-learning point of view, it is common to try to solve a task by fitting a function to sensor or observation data, and thus to try to achieve good performance on the same (or a similar) task when new data is available. The data could, for example, be streaming images from a camera mounted on an indoor mobile robotic platform, and the supervised machine-learning module could learn when and how to turn the wheels left and right, based on collected and labeled visual feeds in a similar indoor environment. Throughout this chapter we have treated the concepts of soft morphology with the repercussions of what are known as morphological processing, sensorimotor coordinated behavior, and soft environment interactions. In similar cases to the example above, it is common for this interconnection of mind, body, and environment to be neglected. In fact, in soft robotics,

as well as other robotics areas, the data is usually perceptual information collected by the robot itself. The perceptual information here is influenced by the morphology of the robot's body, as well as the way in which the robot interacts with entities in the world. The soft robot can thus be seen as a reality filter, which can act in its environment and affect the information in the way most appropriate for learning.

Previous research has shown robots to be capable of purposefully affecting the information gathered from their environment through both morphological processing and sensorimotor coordination (Pfeifer and Scheier 1997; Pfeifer, Iida, and Gómez 2006). In this context, not only the information can be structured so it is rendered suitable for learning, but the structure information itself can guide both the morphology and the control of the robot, creating a sensorimotor and morphological adaptation loop capable of intrinsically driving the robot's behavior. We use the term "soft morphological computation" to describe the ability of a soft robot to understand how its own body and actions filter the information retrieved from the world, and change its configuration and interactions accordingly to optimize information retrieval. This simplification can then drive learning and further the adaptive capabilities of autonomous robotics systems. In Scimeca, Maiolino, and Iida (2018), for example, the soft morphology of the robot is shown to be capable of achieving the cluster separation of stimuli belonging to different object types. Learning can therefore be achieved with unsupervised methods, as the "labels" or classes come from skillful body-environment interaction, which induces sensory separation.

The ability of robotics systems to purposefully shape the sensory information through their actions, or morphology, and to learn from the induced structure has the potential to change the learning landscape within robotics systems. In this context, learning may be thought of not as a process that starts in the information world but rather as one that exists in the physical world, where "learning" the actions and interactions appropriate for sensory perception is the first step toward appropriate learning of the sensory stimuli at a later stage.

Additional Reading and Resources

- A comprehensive review of papers on soft robotics (up to 2007): Trivedi, Deepak, Christopher D. Rahn, William M. Kier, and Ian D. Walker. 2008. "Soft Robotics: Biological Inspiration, State of the Art, and Future Research." *Applied Bionics and Biomechanics* 5 (3): 99–117.
- Paper extensively discussing the connection between cognition, body morphology, and material properties: Pfeifer, Rolf, Fumiya Iida, and Max Lungarella. 2014. "Cognition from the Bottom Up: On Biological Inspiration, Body Morphology, and Soft Materials." *Trends in Cognitive Sciences* 18 (8): 404–413.
- Recent overview of current research, technologies, and applications of soft robotics: Laschi, Cecilia, Jonathan Rossiter, Fumiya Iida, Matteo Cianchetti, and Laura Margheri. *Soft Robotics: Trends, Applications and Challenges. Proceedings of the Soft Robotics Week*. Berlin: Springer.
- Soft robotic tool kit website: <https://softroboticstoolkit.com>.
- Soft robotics TC website: <http://softrobotics.org>.

References

- Asada, Minoru, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. 2009. "Cognitive Developmental Robotics: A Survey." *IEEE Transactions on Autonomous Mental Development* 1 (1): 12–34.
- Asada, Minoru, Karl F. MacDorman, Hiroshi Ishiguro, and Yasuo Kuniyoshi. 2001. "Cognitive Developmental Robotics as a New Paradigm for the Design of Humanoid Robots." *Robotics and Autonomous Systems* 37 (2–3): 185–193.
- Baldassarre, Gianluca, and Marco Mirolli, eds. 2013. *Intrinsically Motivated Learning in Natural and Artificial Systems*. Berlin: Springer.
- Barto, Andrew G. 2013. "Intrinsic Motivation and Reinforcement Learning." In *Intrinsically Motivated Learning in Natural and Artificial Systems*, 17–47. Berlin: Springer.
- Beal, D. N., F. S. Hover, M. S. Triantafyllou, J. C. Liao, and George V. Lauder. 2006. "Passive Propulsion in Vortex Wakes." *Journal of Fluid Mechanics* 549:385–402.
- Belanger, Jim H., and Barry A. Trimmer. 2000. "Combined Kinematic and Electromyographic Analyses of Proleg Function during Crawling by the Caterpillar *Manduca Sexta*." *Journal of Comparative Physiology A* 186 (11): 1031–1039.
- Bernshtein, N. A. 1967. *The Co-ordination and Regulation of Movements*. Oxford: Pergamon Press.
- Berthouze, Luc, and Max Lungarella. 2004. "Motor Skill Acquisition under Environmental Perturbations: On the Necessity of Alternate Freezing and Freeing of Degrees of Freedom." *Adaptive Behavior* 12 (1): 47–64.
- Braitenberg, Valentino. 1986. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
- Brodbeck, Luzius, Simon Hauser, and Fumiya Iida. 2015. "Morphological Evolution of Physical Robots through Model-Free Phenotype Development." *PLoS One* 10 (6): e0128444.
- Brooks, Rodney A. 1990. "Elephants Don't Play Chess." *Robotics and Autonomous Systems* 6 (1–2): 3–15.
- Brown, Eric, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R. Zakin, Hod Lipson, and Heinrich M. Jaeger. 2010. "Universal Robotic Gripper Based on the Jamming of Granular Material." *Proceedings of the National Academy of Sciences* 107 (44): 18809–18814.
- Buhrmann, Thomas, Ezequiel Alejandro Di Paolo, and Xabier Barandiaran. 2013. "A Dynamical Systems Account of Sensorimotor Contingencies." *Frontiers in Psychology* 4:285.
- Cheng, Shi, and Zhigang Wu. 2012. "Microfluidic Electronics." *Lab on a Chip* 12 (16): 2782–2791.
- Cianchetti, M., A. Arienti, M. Follador, B. Mazzolai, P. Dario, and C. Laschi. 2011. "Design Concept and Validation of a Robotic Arm Inspired by the Octopus." *Materials Science and Engineering C* 31 (6): 1230–1239.
- Cianchetti, Matteo, Tommaso Ranzani, Giada Gerboni, Iris De Falco, Cecilia Laschi, and Arianna Menciassi. 2013. "STIFF-FLOP Surgical Manipulator: Mechanical Design and Experimental Characterization of the Single Module." In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3576–3581. New York: IEEE.
- Clark, Andy, and Rick Grush. 1999. "Towards a Cognitive Robotics." *Adaptive Behavior* 7 (1): 5–16.
- Collins, Steve, Andy Ruina, Russ Tedrake, and Martijn Wisse. 2005. "Efficient Bipedal Robots Based on Passive-Dynamic Walkers." *Science* 307 (5712): 1082–1085.
- Dawkins, Richard. 1982. *The Extended Phenotype*. Vol. 8. Oxford: Oxford University Press.
- Della Santina, Cosimo, Robert K. Katzschmann, Antonio Biechi, and Daniela Rus. 2018. "Dynamic Control of Soft Robots Interacting with the Environment." In *2018 IEEE International Conference on Soft Robotics*, 46–53. New York: IEEE.
- de Vignemont, Frédérique. 2010. "Body Schema and Body Image—Pros and Cons." *Neuropsychologia* 48 (3): 669–680.
- Doncieux, Stephane, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E. Gusz Eiben. 2015. "Evolutionary Robotics: What, Why, and Where To." *Frontiers in Robotics and AI* 2:4.
- Edelman, Gerald M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.
- Eder, M., Florian Hisch, and Helmut Hauser. 2018. "Morphological Computation-Based Control of a Modular, Pneumatically Driven, Soft Robotic Arm." *Advanced Robotics* 32 (7): 375–385.
- Fishel, Jeremy A., and Gerald E. Loeb. 2012. "Bayesian Exploration for Intelligent Identification of Textures." *Frontiers in Neurobotics* 6:4.
- Fleming, Peter J., and Robin C. Purshouse. 2002. "Evolutionary Algorithms in Control Systems Engineering: A Survey." *Control Engineering Practice* 10 (11): 1223–1241.

- Floreano, Dario, Ramon Pericet-Camara, Stéphane Viollet, Franck Ruffier, Andreas Brückner, Robert Leitel, Wolfgang Buss et al. 2013. "Miniature Curved Artificial Compound Eyes." *Proceedings of the National Academy of Sciences* 110 (23): 9267–9272.
- Fodor, Jerry A. 1981. *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Fogel, Alan. 2011. "Theoretical and Applied Dynamic Systems Research in Developmental Science." *Child Development Perspectives* 5 (4): 267–272.
- Galloway, Kevin C., Yue Chen, Emily Templeton, Brian Rife, Isuru S. Godage, and Eric J. Barth. 2019. "Fiber Optic Shape Sensing for Soft Robotics." *Soft Robotics* 6 (5): 671–684.
- Geng, Shineng, Youyu Wang, Cong Wang, and Rongjie Kang. 2018. "A Space Tendon-Driven Continuum Robot." In *International Conference on Sensing and Imaging*, 25–35. Cham, Switzerland: Springer.
- Goldfield, Eugene Curtis. 1995. *Emergent Forms: Origins and Early Development of Human Action and Perception*. Oxford: Oxford University Press on Demand.
- Grush, Rick. 2003. "In Defense of Some 'Cartesian' Assumptions Concerning the Brain and Its Operation." *Biology and Philosophy* 18 (1): 53–93.
- Hall, Brian, and Monroe W. Strickberger. 2008. *Strickberger's Evolution*. Burlington, MA: Jones and Bartlett.
- Hara, Fumio, and Rolf Pfeifer, eds. 2003. *Morpho-functional Machines: The New Species: Designing Embodied Intelligence*. Berlin: Springer Science and Business Media.
- Hawkes, Elliot W., Laura H. Blumenschein, Joseph D. Greer, and Allison M. Okamura. 2017. "A Soft Robot That Navigates Its Environment through Growth." *Science Robotics* 2 (8): eaan3028.
- Hoffmann, Matej, Hugo Marques, Alejandro Arieta, Hidenobu Sumioka, Max Lungarella, and Rolf Pfeifer. 2010. "Body Schema in Robotics: A Review." *IEEE Transactions on Autonomous Mental Development* 2 (4): 304–324.
- Hoffmann, Matej, Zdeněk Straka, Igor Farkaš, Michal Vavrečka, and Giorgio Metta. 2017. "Robotic Homunculus: Learning of Artificial Skin Representation in a Humanoid Robot Motivated by Primary Somatosensory Cortex." *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 163–176.
- Homberg, Bianca S., Robert K. Katzschmann, Mehmet R. Dogar, and Daniela Rus. 2015. "Haptic Identification of Objects Using a Modular Soft Robotic Gripper." In *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1698–1705. New York: IEEE.
- Hua, Qilin, Junlu Sun, Haitao Liu, Rongrong Bao, Ruomeng Yu, Junyi Zhai, Caofeng Pan, and Zhong Lin Wang. 2018. "Skin-Inspired Highly Stretchable and Conformable Matrix Networks for Multifunctional Sensing." *Nature Communications* 9 (1): 1–11.
- Hughes, J. A. E., P. Maiolino, and Fumiya Iida. 2018. "An Anthropomorphic Soft Skeleton Hand Exploiting Conditional Models for Piano Playing." *Science Robotics* 3 (25): eaau3098.
- Hughes, Josie, Utku Culha, Fabio Giardina, Fabian Guenther, Andre Rosendo, and Fumiya Iida. 2016. "Soft Manipulators and Grippers: A Review." *Frontiers in Robotics and AI* 3:69.
- Hughes, Josie, and Fumiya Iida. 2017. "Localized Differential Sensing of Soft Deformable Surfaces." In *2017 IEEE International Conference on Robotics and Automation*, 4959–4964. New York: IEEE.
- Iida, Fumiya, and Cecilia Laschi. 2011. "Soft Robotics: Challenges and Perspectives." *Procedia Computer Science* 7:99–102.
- Iida, Fumiya, and Surya G. Nurzaman. 2016. "Adaptation of Sensor Morphology: An Integrative View of Perception from Biologically Inspired Robotics Perspective." *Interface Focus* 6 (4): 20160016.
- Iida, Fumiya, and Rolf Pfeifer. 2006. "Sensing through Body Dynamics." *Robotics and Autonomous Systems* 54 (8): 631–640.
- Jakobi, Nick, Phil Husbands, and Inman Harvey. 1995. "Noise and the Reality Gap: The Use of Simulation in Evolutionary Robotics." In *European Conference on Artificial Life*, 704–720. Berlin: Springer.
- Katzschmann, Robert K., Joseph DelPreto, Robert MacCurdy, and Daniela Rus. 2018. "Exploration of Underwater Life with an Acoustically Controlled Soft Robotic Fish." *Science Robotics* 3 (16): eaar3449.
- Kim, Sangbae, Cecilia Laschi, and Barry Trimmer. 2013. "Soft Robotics: A Bioinspired Evolution in Robotics." *Trends in Biotechnology* 31 (5): 287–294.
- Kuhl, Patricia K. 2000. "Language, Mind, and Brain: Experience Alters Perception." *New Cognitive Neurosciences* 2:99–115.
- Land, Michael F., and Dan-Eric Nilsson. 2012. *Animal Eyes*. Oxford: Oxford University Press.
- Laschi, Cecilia, and Matteo Cianchetti. 2014. "Soft Robotics: New Perspectives for Robot Bodyware and Control." *Frontiers in Bioengineering and Biotechnology* 2:3.

- Laschi, Cecilia, Matteo Cianchetti, Barbara Mazzolai, Laura Margheri, Maurizio Follador, and Paolo Dario. 2012. "Soft Robot Arm Inspired by the Octopus." *Advanced Robotics* 26 (7): 709–727.
- Laschi, Cecilia, Jonathan Rossiter, Fumiya Iida, Matteo Cianchetti, and Laura Margheri. 2016. *Soft Robotics: Trends, Applications and Challenges*. Berlin: Springer.
- Li, S. H., Q. Y. Zeng, Y. L. Xiao, S. Y. Fu, and B. L. Zhou. 1995. "Biomimicry of Bamboo Bast Fiber with Engineering Composite Materials." *Materials Science and Engineering C* 3 (2): 125–130.
- Lin, Huai-Ti, Gary G. Leisk, and Barry Trimmer. 2011. "GoQBot: A Caterpillar-Inspired Soft-Bodied Rolling Robot." *Bioinspiration and Biomimetics* 6 (2): 026007.
- Lipson, Hod, and Jordan B. Pollack. 2000. "Automatic Design and Manufacture of Robotic Lifeforms." *Nature* 406 (6799): 974–978.
- Lu, Nanshu, and Dae-Hyeong Kim. 2014. "Flexible and Stretchable Electronics Paving the Way for Soft Robotics." *Soft Robotics* 1 (1): 53–62.
- Lund, Henrik Hautop, John Hallam, and Wei-Po Lee. 1997. "Evolving Robot Morphology." In *Proceedings of the 1997 IEEE International Conference on Evolutionary Computation*, 197–202. New York: IEEE.
- Lungarella, Max, Giorgio Metta, Rolf Pfeifer, and Giulio Sandini. 2003. "Developmental Robotics: A Survey." *Connection Science* 15 (4): 151–190.
- Maiolino, P., Fabia Galantini, F. Mastrogiovanni, G. Gallone, G. Cannata, and Federico Carpi. 2015. "Soft Dielectrics for Capacitive Sensing in Robot Skins: Performance of Different Elastomer Types." *Sensors and Actuators A: Physical* 226:37–47.
- Maiolino, Perla, Marco Maggiali, Giorgio Cannata, Giorgio Metta, and Lorenzo Natale. 2013. "A Flexible and Robust Large Scale Capacitive Tactile System for Robots." *IEEE Sensors Journal* 13 (10): 3910–3917.
- Majidi, Carmel. 2014. "Soft Robotics: A Perspective—Current Trends and Prospects for the Future." *Soft Robotics* 1 (1): 5–11.
- Matarić, M. J., and B. Scassellati. 2016. "Socially Assistive Robotics." In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1973–1994. Berlin: Springer.
- Mautner, Craig, and Richard K. Belew. 2000. "Evolving Robot Morphology and Control." *Artificial Life and Robotics* 4 (3): 130–136.
- Mazzolai, Barbara, Lucia Beccai, and Virgilio Mattoli. 2014. "Plants as Model in Biomimetics and Biorobotics: New Perspectives." *Frontiers in Bioengineering and Biotechnology* 2:2.
- Mazzolai, Barbara, Alessio Mondini, Paolo Corradi, Cecilia Laschi, Virgilio Mattoli, Edoardo Sinibaldi, and Paolo Dario. 2010. "A Miniaturized Mechatronic System Inspired by Plant Roots for Soil Exploration." *IEEE/ASME Transactions on Mechatronics* 16 (2): 201–212.
- Miall, R. C., D. Jo Weir, Daniel M. Wolpert, and J. F. Stein. 1993. "Is the Cerebellum a Smith Predictor?" *Journal of Motor Behavior* 25 (3): 203–216.
- Miriyev, Aslan, Kenneth Stack, and Hod Lipson. 2017. "Soft Material for Soft Actuators." *Nature Communications* 8 (1): 1–8.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, and S. Petersen. 2015. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518 (7540): 529–533.
- Muth, Joseph T., Daniel M. Vogt, Ryan L. Truby, Yiğit Mengüç, David B. Kolesky, Robert J. Wood, and Jennifer A. Lewis. 2014. "Embedded 3D Printing of Strain Sensors within Highly Stretchable Elastomers." *Advanced Materials* 26 (36): 6307–6312.
- Nabeshima, Cota, Yasuo Kuniyoshi, and Max Lungarella. 2006. "Adaptive Body Schema for Robotic Tool-Use." *Advanced Robotics* 20 (10): 1105–1126.
- Nakajima, Kohei, Helmut Hauser, Rongjie Kang, Emanuele Guglielmino, Darwin G. Caldwell, and Rolf Pfeifer. 2013. "A Soft Body as a Reservoir: Case Studies in a Dynamic Model of Octopus-Inspired Soft Robotic Arm." *Frontiers in Computational Neuroscience* 7:91.
- Nakajima, Kohei, Helmut Hauser, Tao Li, and Rolf Pfeifer. 2015. "Information Processing via Physical Soft Body." *Scientific Reports* 5:10487.
- Nakajima, Kohei, Tao Li, Helmut Hauser, and Rolf Pfeifer. 2014. "Exploiting Short-Term Memory in Soft Body Dynamics as a Computational Resource." *Journal of the Royal Society Interface* 11 (100): 20140437.
- Newell, Allen, and Herbert A. Simon. 2007. "Computer Science as Empirical Inquiry: Symbols and Search." In *ACM Turing Award Lectures*, 1975. New York: Association for Computing Machinery.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.

- Nolfi, Stefano, Dario Floreano, Orazio Miglino, and Francesco Mondada. 1994. "How to Evolve Autonomous Robots: Different Approaches in Evolutionary Robotics." In Vol. 4, *Artificial Life: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, 190–197. Cambridge, MA: MIT Press.
- O'Halloran, Ailish, Fergal O'Malley, and Peter McHugh. 2008. "A Review on Dielectric Elastomer Actuators, Technology, Applications, and Challenges." *Journal of Applied Physics* 104 (7): 9.
- O'Regan, J. Kevin, and Alva Noë. 2001. "A Sensorimotor Account of Vision and Visual Consciousness." *Behavioral and Brain Sciences* 24 (5): 939–973.
- Oudeyer, Pierre-Yves, Frédéric Kaplan, and Verena V. Hafner. 2007. "Intrinsic Motivation Systems for Autonomous Mental Development." *IEEE Transactions on Evolutionary Computation* 11 (2): 265–286.
- Parker, Sue Taylor, and Michael L. McKinney. 2012. *Origins of Intelligence: The Evolution of Cognitive Development in Monkeys, Apes, and Humans*. Baltimore: Johns Hopkins University Press.
- Pearson, Martin J., Ben Mitchinson, J. Charles Sullivan, Anthony G. Pipe, and Tony J. Prescott. 2011. "Biomimetic Vibrissal Sensing for Robots." *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1581): 3085–3096.
- Pfeifer, Rolf. 2000. "On the Role of Morphology and Materials in Adaptive Behavior." *From Animals to Animals* 6:23–32. Cambridge, MA: MIT Press.
- Pfeifer, Rolf, Fumiya Iida, and Josh Bongard. 2005. "New Robotics: Design Principles for Intelligent Systems." *Artificial Life* 11 (1–2): 99–120.
- Pfeifer, Rolf, Fumiya Iida, and Gabriel Gómez. 2006. "Morphological Computation for Adaptive Behavior and Cognition." In Vol. 1291, *International Congress Series*, 22–29. San Diego: Elsevier.
- Pfeifer, Rolf, Fumiya Iida, and Max Lungarella. 2014. "Cognition from the Bottom Up: On Biological Inspiration, Body Morphology, and Soft Materials." *Trends in Cognitive Sciences* 18 (8): 404–413.
- Pfeifer, Rolf, Max Lungarella, and Fumiya Iida. 2007. "Self-Organization, Embodiment, and Biologically Inspired Robotics." *Science* 318 (5853): 1088–1093.
- Pfeifer, Rolf, Max Lungarella, and Fumiya Iida. 2012. "The Challenges Ahead for Bio-inspired 'Soft' Robotics." *Communications of the ACM* 55 (11): 76–87.
- Pfeifer, Rolf, and Christian Scheier. 1997. "Sensory-Motor Coordination: The Metaphor and Beyond." *Robotics and Autonomous Systems* 20 (2–4): 157–178.
- Pfeiffer, Rolf, and Christian Scheier. 1999. *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Piaget, Jean. 2003. *The Psychology of Intelligence*. London: Routledge.
- Piaget, Jean, and Margaret Cook. 1952. *The Origins of Intelligence in Children*. Vol. 8. New York: International Universities Press.
- Prescott, Tony J., Martin J. Pearson, Ben Mitchinson, J. Charles W. Sullivan, and Anthony G. Pipe. 2009. "Whisking with Robots." *IEEE Robotics and Automation Magazine* 16 (3): 42–50.
- Rich, Steven I., Robert J. Wood, and Carmel Majidi. 2018. "Untethered Soft Robotics." *Nature Electronics* 1 (2): 102–112.
- Rochat, Philippe. 1998. "Self-Perception and Action in Infancy." *Experimental Brain Research* 123 (1–2): 102–109.
- Rodrigue, Hugo, Wei Wang, Min-Woo Han, Thomas J. Y. Kim, and Sung-Hoon Ahn. 2017. "An Overview of Shape Memory Alloy-Coupled Actuators and Robots." *Soft Robotics* 4 (1): 3–15.
- Rodrigue, Hugo, Wei Wang, Dong-Ryul Kim, and Sung-Hoon Ahn. 2017. "Curved Shape Memory Alloy-Based Soft Actuators and Application to Soft Gripper." *Composite Structures* 176:398–406.
- Rogers, John A., Takao Someya, and Yonggang Huang. 2010. "Materials and Mechanics for Stretchable Electronics." *Science* 327 (5973): 1603–1607.
- Rosendo, Andre, Marco von Atzigen, and Fumiya Iida. 2017. "The Trade-Off between Morphology and Control in the Co-optimized Design of Robots." *PloS One* 12 (10): e0186107.
- Rucker, D. Caleb, and Robert J. Webster. 2014. "Mechanics of Continuum Robots with External Loading and General Tendon Routing." In *Experimental Robotics*, edited by Oussama Khatib, Vijay Kumar, and Gaurav Sukhatme, 645–654. Berlin: Springer.
- Rus, Daniela, and Michael T. Tolley. 2015. "Design, Fabrication and Control of Soft Robots." *Nature* 521 (7553): 467–475.
- Sadeghi, Ali, Alice Tonazzini, Liyana Popova, and Barbara Mazzolai. 2013. "Robotic Mechanism for Soil Penetration Inspired by Plant Root." In *2013 IEEE International Conference on Robotics and Automation*, 3457–3462. New York: IEEE.

- Sadeghi, Ali, Alice Tonazzini, Liyana Popova, and Barbara Mazzolai. 2014. "A Novel Growing Device Inspired by Plant Root Soil Penetration Behaviors." *PLoS One* 9 (2): e90139.
- Scheier, Christian, and Rolf Pfeifer. 1999. "The Embodied Cognitive Science Approach." In *Dynamics, Synergetics, Autonomous Agents: Nonlinear Systems Approaches to Cognitive Psychology and Cognitive Science*, edited by W. Tschacher and J.-P. Dauwalder, 159–179. Singapore: World Scientific.
- Scheltjens, G., M. M. Diaz, J. Brancart, G. Van Assche, and B. Van Mele. 2013. "A Self-Healing Polymer Network Based on Reversible Covalent Bonding." *Reactive and Functional Polymers* 73 (2): 413–420.
- Schmidhuber, Jürgen. 2015. "Deep Learning in Neural Networks: An Overview." *Neural Networks* 61:85–117.
- Scimeca, Luca, Josie Hughes, Perla Maiolino, and Fumiya Iida. 2019. "Model-Free Soft-Structure Reconstruction for Proprioception Using Tactile Arrays." *IEEE Robotics and Automation Letters* 4 (3): 2479–2484.
- Scimeca, Luca, Perla Maiolino, and Fumiya Iida. 2018. "Soft Morphological Processing of Tactile Stimuli for Autonomous Category Formation." In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, 356–361. New York: IEEE.
- Scimeca, Luca, Perla Maiolino, and Fumiya Iida. 2020. "Efficient Bayesian Exploration for Soft Morphology-Action Co-optimization." In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 639–644. New York: IEEE.
- Scott-Phillips, Thomas C., Kevin N. Laland, David M. Shuker, Thomas E. Dickins, and Stuart A. West. 2014. "The Niche Construction Perspective: A Critical Appraisal." *Evolution* 68 (5): 1231–1243.
- Seok, Sangok, Cagdas D. Onal, Robert Wood, Daniela Rus, and Sangbae Kim. 2010. "Peristaltic Locomotion with Antagonistic Actuators in Soft Robotics." In *2010 IEEE International Conference on Robotics and Automation*, 1228–1233. New York: IEEE.
- Shintake, Jun, Bryan Schubert, Samuel Rosset, Herbert Shea, and Dario Floreano. 2015. "Variable Stiffness Actuator for Soft Robotics Using Dielectric Elastomer and Low-Melting-Point Alloy." In *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1097–1102. New York: IEEE.
- Shumaker, Robert W., Kristina R. Walkup, and Benjamin B. Beck. 2011. *Animal Tool Behavior: The Use and Manufacture of Tools by Animals*. Baltimore: Johns Hopkins University Press.
- Siegenthaler, K. O., A. Künkel, G. Skupin, and M. Yamamoto. 2011. "Ecoflex® and Ecovio®: Biodegradable, Performance-Enabling Plastics." In *Synthetic Biodegradable Polymers*, edited by Bernhard Rieger, Andreas Kunkel, Geoffrey W. Coates, Robert Reichardt, Eckhard Dinjus, and Thomas A. Zevaco, 91–136. Berlin: Springer.
- Soni, Mahesh, and Ravinder Dahiya. 2020. "Soft eSkin: Distributed Touch Sensing with Harmonized Energy and Computing." *Philosophical Transactions of the Royal Society A* 378 (2164): 20190156.
- Sporns, Olaf. 2003. "Embodied Cognition." In *Handbook of Brain Theory and Neural Networks*, edited by Michael A. Arbib. Cambridge, MA: MIT Press.
- Taccola, S., A. Zucca, F. Greco, B. Mazzolai, and V. Mattoli. 2013. "Electrically Driven Dry State Actuators Based on PEDOT: PSS Nanofilms." In *EuroEAP 2013 International Conference on Electromechanically Active Polymer (EAP) Transducers and Artificial Muscles*. Duebendorf, Switzerland, June 25–26.
- Taga, Gentaro, Rieko Takaya, and Yukuo Konishi. 1999. "Analysis of General Movements of Infants towards Understanding of Developmental Principle for Motor Control." In *1999 IEEE International Conference on Systems, Man, and Cybernetics*. Cat. No. 99CH37028. Vol. 5, 678–683. New York: IEEE.
- Terryn, Seppe, Joost Brancart, Dirk Lefeber, Guy Van Assche, and Bram Vanderborght. 2017. "Self-Healing Soft Pneumatic Robots." *Science Robotics* 2 (9): 1–12.
- Thelen, Esther, and Linda B. Smith. 1996. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge: MIT Press.
- Tolley, Michael T., Robert F. Shepherd, Michael Karpelson, Nicholas W. Bartlett, Kevin C. Galloway, Michael Wehner, Rui Nunes, George M. Whitesides, and Robert J. Wood. 2014. "An Untethered Jumping Soft Robot." In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 561–566. New York: IEEE.
- Towal, R. Blythe, Brian W. Quist, Venkatesh Gopal, Joseph H. Solomon, and Mitra J. Z. Hartmann. 2011. "The Morphology of the Rat Vibrissal Array: A Model for Quantifying Spatiotemporal Patterns of Whisker-Object Contact." *PLoS Computational Biology* 7 (4): e1001120.
- Trivedi, Deepak, Christopher D. Rahn, William M. Kier, and Ian D. Walker. 2008. "Soft Robotics: Biological Inspiration, State of the Art, and Future Research." *Applied Bionics and Biomechanics* 5 (3): 99–117.
- Trueman, Edwin Royden. 1975. *Locomotion of Soft-Bodied Animals*. London: Edward Arnold.
- Tsakiris, Manos. 2010. "My Body in the Brain: A Neurocognitive Model of Body-Ownership." *Neuropsychologia* 48 (3): 703–712.

- Tuffield, Paul, and Hugo Elias. 2003. "The Shadow Robot Mimics Human Actions." *Industrial Robot: An International Journal* 30 (1).
- Turvey, Michael T. 1990. "Coordination." *American Psychologist* 45 (8): 938.
- Velcro SA. 1955. "Improvements in or Relating to a Method and a Device for Producing a Velvet Type Fabric." Swiss Patent 721338.
- Vujovic, Vuk, Andre Rosendo, Luzius Brodbeck, and Fumiya Iida. 2017. "Evolutionary Developmental Robotics: Improving Morphology and Control of Physical Robots." *Artificial Life* 23 (2): 169–185.
- Wang, Liyu, Luzius Brodbeck, and Fumiya Iida. 2014. "Mechanics and Energetics in Tool Manufacture and Use: A Synthetic Approach." *Journal of the Royal Society Interface* 11 (100): 20140827.
- Wolpert, Daniel M., Kenji Doya, and Mitsuo Kawato. 2003. "A Unifying Computational Framework for Motor Control and Social Interaction." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 358 (1431): 593–602.
- Yap, Hong Kai, Hui Yong Ng, and Chen-Hua Yeow. 2016. "High-Force Soft Printable Pneumatics for Soft Robotic Applications." *Soft Robotics* 3 (3): 144–158.
- Yirmibesoglu, Osman Dogan, John Morrow, Steph Walker, Walker Gosrich, Reece Cañizares, Hansung Kim, Uranbileg Daalkhajiv, Chloe Fleming, Callie Branyan, and Yigit Menguc. 2018. "Direct 3D Printing of Silicone Elastomer Soft Robots and Their Performance Comparison with Molded Counterparts." In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, 295–302. New York: IEEE.
- Zelik, Karl E., and Arthur D. Kuo. 2010. "Human Walking Isn't All Hard Work: Evidence of Soft Tissue Contributions to Energy Dissipation and Return." *Journal of Experimental Biology* 213 (24): 4257–4264.
- Zlatev, Jordan, and Christian Balkenius. 2001. "Why Epigenetic Robotics." In *First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 1–4. Lund, Sweden: Lund University Cognitive Studies.

II METHODS AND CONCEPTS

7 Robot Platforms and Simulators

Diego Ferigo, Alberto Parmiggiani, Elena Rampone, Vadim Tikhanoff,
Silvio Traversaro, Daniele Pucci, and Lorenzo Natale

7.1 Introduction

Cognitive robotics is a broad field that spans diverse areas of robotics such as human-robot interaction (HRI), navigation, visual perception, object manipulation, physical human-robot interaction, and the study of cognitive architectures. This places specific constraints on the robotic platform to be used. HRI, for example, studies robot behaviors that are as close as possible to those of humans, with the goal of making the interaction between robots and humans as seamless as possible. HRI relies on communication channels that are familiar to humans, such as speech, vision, and touch. To implement humanlike robot behaviors some HRI studies require a platform capable of replicating at least some of the movements of humans (such as eye movements or gestures). Navigation and visual perception are typically carried out using a combination of sensors, such as LIDAR, RGB, or RGBD cameras. Object manipulation and physical human-robot interaction benefit from torque sensors and tactile sensors. The study of cognitive architectures is often bioinspired; it emphasizes humanlike sensing and perception and, often, their integration in multimodal studies.

For these reasons, the focus of research in cognitive robotics is frequently on systems-level capabilities. In these cases, individual capabilities are not to be studied in isolation and must be integrated into the same platform.

Finally, given its intrinsic interdisciplinary nature, research in cognitive robotics is carried out not only by roboticists but also by computer scientists, psychologists, and neuroscientists with little expertise in mechatronics.

It is not surprising, therefore, that the community of researchers working in cognitive robotics has been among the first to recognize the importance of the platform as an enabler in investigating given research questions and, in addition, to highlight the advantages of research platforms that are easy to use by nonexperts and that are shared among different groups. Early examples of platforms adopted in the cognitive robotic community are the Aibo Robot developed by the Sony Corporation (2020) and the iCub humanoid robot (Parmiggiani et al. 2012). Other popular examples are the NAO (Gouaillier et al. 2009) and Pepper (Pandey and Gelin 2018) robots developed by Aldebaran.

A large amount of research in robotics is carried out in simulation. This is because software simulators allow much faster prototyping and debugging, especially considering

that most robotic platforms are prototypes with limited reliability. Developing software in simulation allows for testing research algorithms without the worries of damaging the robot or the environment. Recently, deep-learning research has demonstrated that it is possible to train algorithms using data generated in simulation and deploy them in the real world with impressive results. This has been shown to work well in perception using a mix of data augmentation and photorealistic rendering to solve the problem of six-dimensional object pose estimation (Tremblay et al. 2018) and, in reinforcement learning, to solve in-hand object manipulation with a dexterous hand (Andrychowicz et al. 2020). This research has pushed the development of simulation tools that are able to reproduce the physical environment with great accuracy, including sophisticated photorealism.

In the past there have been efforts to develop platforms specifically tailored to research cognitive robots. A notable example is the iCub, a humanoid robot specifically developed to target the cognitive robotic community. Other platforms were not designed with this goal in mind but have become de facto standards thanks to their massive adoption (e.g., Aibo, NAO, and Pepper, already mentioned above, and also Baxter from Rethink Robotics [Fitzgerald 2013] and Panda from Franka Emika GmbH [2020]). The goal of this chapter is to identify and describe the robotic platforms and simulation tools used most often by the community, highlighting their pros and cons in supporting research activities.

7.2 Methodology

In writing this chapter, we tried to understand which platforms are in use in the cognitive robotic community. We performed detailed research by looking at two of the main scientific journals on cognitive robotics: *IEEE Transactions on Cognitive Developmental Systems (TCDS)* and the Springer *International Journal of Social Robotics (IJSR)*. We inspected all papers published in these journals during the period 2016–2019, noting for each which robot platforms and which software simulators (if any) were employed. The goal of this literature survey was to identify those platforms and software simulators used within the community. It is worth mentioning that we decided to focus on journals instead of a larger pool of venues, including conferences, because this allowed us to inspect publications over a longer time span and to have access to more consolidated work.

The results are summarized in table 7.1 for the hardware platforms and table 7.2 for the simulators. In table 7.1 we report, for each platform, the number of times a paper was published either in *TCDS* or *IJSR* during the considered period. Because it was not possible to list all the robot platforms, we grouped all the platforms found in a small number of papers (not more than two times) in the category “others.”

Overall we analyzed 337 papers, in which we found references to 62 different hardware platforms and 13 software frameworks. The first observation was the large fragmentation of the community: with minor exceptions, most robotic platforms were found only once. This demonstrates that other groups do not use these platforms and that most research in the field is carried out with custom prototypes used only within a specific research group. Yet this investigation also allowed us to clearly identify some platforms that are used within the community: The NAO robot was found to be used most often (forty-seven papers), followed by iCub (sixteen), Pepper and Robovie (seven), the Pioneer 3-Dx/3-AT (six), Baxter (five),

Table 7.1
Literature survey: results for the hardware platforms

	<i>TCDS</i> 2019	<i>TCDS</i> 2018	<i>TCDS</i> 2017	<i>TCDS</i> 2016	<i>IJSR</i> 2019	<i>IJSR</i> 2018	<i>IJSR</i> 2017	<i>IJSR</i> 2016	SUM
Others	10	9	3	4		10	16	10	62
NAO	3	5	3	2	4	14	9	7	47
iCub		6	2	6			2		16
Pepper	1				3		3		7
Robovie family						3	2	2	7
Pioneer 3-Dx/3-AT	2	1	2			1			6
Baxter	1	2		1			1		5
Kuka LBR iiwa				1	1		1		3

Note: We report the number of papers that were published on each platform for each year in the *TCDS* and *IJSR*. The last column reports the sum across all years. We list only those platforms found in more than two papers. Overall we found sixty-seven different platforms.

Table 7.2
Literature survey: simulation frameworks

Simulator	Robot platform	Supported language	Operating system	Reference
ODE	iCub (5), HOAP-2 (1)	C/C++	macOS, Linux, Windows	(Smith 2020)
Gazebo	NeuroSnake (1), Eddie (1), PKU-HR6 (1)	C++	macOS, Linux, Windows	(Open Software Robotics Foundation 2014)
V-REP	NAO (1), Pioneer (2), Exapod (1)	LUA, C/C++, Python, Java, MATLAB/Octave	macOS, Linux, Windows	(Coppelia Robotics GmbH 2020)
Webots	NAO (1), Salamander (1)	C/C++, Python, Java, MATLAB	macOS, Linux, Windows	(Cyberbotics Ltd. 2020)
Nextage	Kawada (1)	Python		
Blender	Custom (1)	Python	macOS, Linux, Windows	(Blender Foundation 2020)
OpenSim	JacoArm (1)	C++, Python, MATLAB, Java	macOS, Linux (API only), Windows	(OpenSim 2020)
RobWorkSim	UR3 (1)	C++, Python, Java, LUA	macOS, Linux, Windows	(SDU Robotics 2020)
SMILE	Baxter (1)	Java	macOS, Linux, Windows	(SMILE 2020)
Stage	iRat (1)	C++	macOS, Linux	(Stage 2020)
FARSA	iCub (1)	C++	macOS, Linux, Windows	(Farsa Sourceforge 2020)
SIGVerse	TurtleBot (1)	C# (Unity)	Windows	(SIGVerse 2020)
3D Studio	Probo (1)	macOS, Linux, Windows	Windows	(Autodesk Inc. 2020)

Note: For each simulation framework, we report the name of the robot simulated and the number of times it occurred in the papers we analyzed. We also add general information on the supported programming languages and operating system and link to the web page hosting the simulator code.

and Kuka LBR iiwa (three). These platforms are reviewed in some detail in the next section of this chapter.

We also observed that, surprisingly, software simulators are not used very often in the community, which largely prefers to experiment with real robots. We believe this is because simulators are still quite immature for research in cognitive robotics, as they do not model complex environments well, especially when interaction with humans is important. Another possibility could be that recent progress in software simulation has been mostly pushed by the robot learning community, which focuses on perception and grasping and publishes in different venues (e.g., the International Conference on Robot Learning). The following software frameworks were found to be used most often: Open Dynamic Engine (ODE; four times), Gazebo (three times), and V-REP (four times). However, this view does not represent well the growing importance that the robotic community has given to the development of simulation frameworks. This situation will change, as it is likely that much better simulation environments will be available in the coming years, with consequential impact on the cognitive robotics community. For these reasons, in section 7.4 we provide an overview of how software simulators are used in robotics, describing the different type of simulators used in the community and the current trends in research pushed by the growing field of robot learning.

7.3 Robot Platforms

In this section we review the platforms most commonly used in the cognitive robotics community, as identified by our survey. As a summary the details of each platform are also reported in table 7.3.

The iCub is an open-source humanoid robot developed within the context of the RobotCub project (figure 7.1a). The iCub has fifty-three degrees of freedom (DOF), and it is endowed with a rich sensor suite (stereo cameras, microphones, six-axis force-torque sensors, and whole-body tactile sensors). The cost of the iCub robot is about €250,000. It is controlled with the YARP middleware and custom motion-control libraries (more details on the software architecture of the iCub are described in Natale et al. [2016]). These elements allow the planning and control of complex HRI tasks that also involve physical interaction. Although the iCub has onboard computation and batteries, it is generally operated from a fixed base that does not allow autonomous navigation. The iCub is a versatile platform used to study all aspects of robotics, including sensorimotor learning (Hoffmann et al. 2018; Zambelli and Demiris 2017; Giagkos et al. 2017; Çelikkanat et al. 2016), object learning and tool use (Ribes et al. 2016; Mar, Tikhonoff, and Natale 2017), intrinsically motivated and reinforcement learning (Meola et al. 2016; Santucci, Baldassarre, and Mirolli 2016), HRI (Förster, Saunders, and Nehaniv 2018; Baraglia, Nagai, and Asada 2016; Petit, Fischer, and Demiris 2016), social robotics (Anzalone et al. 2017; Ivaldi et al. 2017), and artificial cognitive architectures (Moulin-Frier et al. 2018).

The Pepper robot is a humanoid robot originally developed by the Aldebaran company, which was later acquired by Softbank Robotics (figure 7.1c). The Pepper robot has seventeen DOF and is equipped with omnidirectional wheels for navigating indoor environments. The Pepper robot was specifically developed for social, nonphysical HRI tasks; its

Table 7.3
Details of hardware platforms

Robot	Type	Locomotion	Estimated cost	Height [m]	Arm length [m]	Weight [kg]	Typical use	Programming languages
NAO	Humanoid	Legs	\$7,500	0.574	0.22	5.4	Festuring, HRI	NaoQi SDK, C++, Python, Java
iCub	Humanoid	Legs	€250,000	1.05	0.36	33	Manipulation, HRI	C++, YARP
Pepper	Humanoid	Wheels (3), holonomic	\$13,100*	1.21	0.4	28	Navigation, HRI	NaoQi SDK, C++, Python, Java
Robovie (R3)	Humanoid	Wheels (2), diff. drive	\$40,000*	1.08	≈0.32	35	HRI	RobovieMaker2, C++
Pioneer (3-AT)	Mobile base	Wheels (4), diff. drive	\$4,195–\$30,000*	0.24	–	9	Navigation, SLAM	C++, ROS
Baxter	Robot arm	N/A	\$25,000*	0.94	0.97	75	Manipulation, physical HRI	C++, Python, ROS
Kuka LBR iiwa	Robot arm	N/A	\$200,000	–	0.8	22.3 [†]	Manipulation, physical HRI	C++
Panda	Robot arm	N/A	\$20,000	–	0.855	18 [†]	Manipulation, physical HRI	C++, ROS

Note: The cost of the platforms is estimated based on information available on the web and in the literature (see text for details and sources).

*Not available for purchase.

[†]Control unit excluded.



Figure 7.1

(a) The iCub robot. *Source:* Courtesy of the Istituto Italiano di Tecnologia. (b) The NAO robot. *Source:* Wikimedia: Ubahnverleih 2016, released with license CC0. (c) The Pepper. *Source:* Wikimedia: Tokumeigakarinooshima 2014, released with license CC0. (d) Robovie R3. *Source:* Courtesy of the Cognitive and Social Robotics Laboratory, Istanbul Technical University.

motors are sufficiently powerful to move the joints but not strong enough to hurt someone through accidental contact. The Pepper robot was designed with a focus on affordability. The cost of this robot has been reported to be about \$440/month for an enterprise model (TechCrunch 2015) and \$13,100 for the 2018 edition of the RobotCub@Home competitor (RobotCub@Home 2018). The Pepper mechanical structure relies heavily on plastic materials, for structural parts as well as bearings. Pepper is equipped with cameras, three-dimensional sensors for visual perception, and microphones for auditory perception, as well as laser, sonar, and infrared sensors for navigation.

Pepper has been used for HRI (Izui and Venture 2020), including robot-assisted therapy (Cao et al. 2019), emotion recognition (Tsiourti et al. 2019), and communication (Hirano et al. 2018; Claret, Venture, and Basañez 2017), as well as robot design studies (Kwak, Kim, and Choi 2017).

The NAO is a small humanoid robot developed by the Aldebaran company, later acquired by Softbank Robotics (figure 7.1b). NAO has twenty-five DOF and was designed to be an affordable, open, and easy-to-handle platform. The cost of a NAO is about \$7,500 (Smashing Robotics 2016). It is 57 cm tall and weighs 4.5 kg. Thanks to its simple and functional design,

NAO became the standard platform for the RoboCup league in 2008. It has been a popular choice for groups working in HRI who want to avoid the experimental complexities related to the use of larger robots.

The NAO robot is primarily used for HRI studies (e.g., Khamassi et al. 2019; Murata et al. 2018; Liu and Zhang 2016; Izui and Venture 2020), including robot-mediated therapy for autistic children (Cao et al. 2019; David et al. 2018) and educational robotics (Chandra, Dillenbourg, and Paiva 2020; Jones and Castellano 2018), but also sensorimotor learning (Wieser and Cheng 2018), imitation learning (Park, Kim, and Nagai 2017), and learning by demonstration of tactile gestures (Pierris and Dahl 2017), to mention just a few.

A considerable amount of research on HRI, especially in Japan, is carried out on the Robovie R2 platform and its successor R3 (ATR-Creative 2010; figure 7.1*d*). The Robovie R3 is a small, 108 cm tall humanoid robot that has two four-DOF arms, a three-DOF neck, and two wheels for autonomous navigation. It carries two video-camera touch sensors and a laser for detecting obstacles. The robot was developed by ATR-C and VStone in 2010 and was supplied until 2016 at the price of about \$40,000 (RobotShop Community 2020), when it was eventually discontinued.

The robots of the Robovie series have been extensively used for researching HRI, assessing anthropomorphism (Złotowski et al. 2018), testing deictic behavior (Liu et al. 2017), testing the effects of negative evaluations (Nomura and Kanda 2015), evaluating lexical entrainment (Iio et al. 2015), and studying social side-by-side walking (Karunaratne et al. 2018).

Other robots without an anthropomorphic appearance were employed frequently in the papers analyzed in our study. Baxter was presented in 2012 by the American company Rethink Robotics. It is a bimanual manipulator with two seven-DOF arms (figure 7.2*a*). The Baxter was developed with a focus on safe physical HRI and was therefore equipped with series-elastic actuators (SEAs) at all arm joints. This feature allows the robot to perceive external forces and consequently adapt its motion. Rethink Robotics designed the Baxter robot to be economically viable, targeting repetitive assembly applications in small and medium enterprises (SMEs). The robot was supplied at the average cost equivalent to the salary of an assembly operator in the United States (base price of \$25,000; Wikipedia 2020). In 2018 Rethink Robotics ceased operations, thus interrupting the Baxter program.

In our survey we found the Baxter robot used to study learning by demonstration (Yang et al. 2018) and learning by imitation (Katz et al. 2018) and to evaluate HRIs (Herath, Jochum, and Vlachos 2018) and assess the legibility of behaviors in collaborative tasks (Busch et al. 2017). The LBR iiwa from Kuka AG is a seven-DOF robot arm, with integrated joint-level torque sensing (figure 7.2*b*). It is based on the hardware of the DLR LWR III (Hirzinger et al. 2002) developed at the DLR Institute for Robotics and Mechatronics and is available for about \$200,000 (*Robotics Business Review* 2015). The capability for torque and force control make this robot especially suitable for experiments that require safe physical interaction.

The Kuka LBR iiwa robot was, for example, used in experiments on cooperative object manipulation (Donner et al. 2017), affordance learning (Ugur and Piater 2017), and learning by demonstration in an assistive context (Laurettili, Cordella, and Zollo 2019). A second robot arm suited for physical HRI experiments is the Panda from Franka Emika GmbH




Figure 7.2

(a) The Baxter robot. *Source:* Energy.gov 2013, released with license CC0. (b) Kuka LBR iiwa. *Source:* Caré 2015, released with license CC BY-SA 4.0. (c) The Panda arm. *Source:* Ims 2017, released with license CC BY-SA 4.0. (d) The Pioneer robot (3-AT model), equipped with a gripper. *Source:* J. Wang 2008, released with license CC BY-SA 3.0.

(figure 7.2c). The Panda has seven DOF; each joint integrates joint torque sensing at 1 kHz. One of the distinctive features of this system is its relatively low price tag: its target cost was €10,000 (*IEEE Spectrum* 2020a), and at the time of writing, in France it is distributed with its software for about €20,000 excluding taxes (Generation Robots 2020). A second advantage of the Panda is its user-friendly programming interface, which makes it accessible to users with no expertise in software programming. Overall, these features make the Panda arm a popular choice for research in collaborative robotics and object grasping. The Pioneer 3-DX/3-AT models by Adept (now part of Omron) are wheeled mobile robots that have been used extensively for research (figure 7.2d). The Pioneer 3-DX is a compact differential-drive mobile robot, with two motorized wheels. It comprises a motion controller and sensors (sonars and optional laser scanner) for navigation and obstacle avoidance. The Pioneer 3-AT is a similar platform with four wheels designed for outdoor navigation. The Pioneer robots were discontinued in 2015 as part of the new Omron product strategy after the acquisition of Adept. The cost of the platform, depending on the configuration, varied between \$4,195 to \$30,000 (*IEEE Spectrum* 2020a). They were often equipped with additional sensors (like RGBD or RGBD cameras) and grippers. Benli, Motai, and Rogers (2019) equipped a 3-DX

with a thermal camera to study human behavior tracking, and Glover and Wyeth (2018) equipped it with a gripper to study the lifelong learning of object affordances. The same platform was used to study how to solve object search tasks by integrating object identification, avoidance, path planning, and navigation (Wang et al. 2019). It was also employed in HRI settings to study attention (Caccavale and Finzi 2017) and to evaluate the effectiveness of telepresence interfaces (Ahn and Kim 2018).

7.4 Software Simulators

In all fields of modern engineering, it is standard practice to automatically build mathematical models that describe systems currently being designed or under study and then to use these models in digital computers to “simulate” their behavior.

One of the advantages of simulations is that they enable research where real-world investigations would be difficult to conduct. For example, the phenomena of interest could be inaccessible, too dangerous, too expensive, or morally unacceptable to study empirically, at least at an early experimental stage. Even though the study of real phenomena is often desirable, simulations provide a set of advantages in comparison to studying the real world. They allow for repeated observations, strict control of conditional parameters, and scalability. In general, simulations offer controlled, safe, and affordable environments in which task-oriented, social, and cognitive skills can be repeatedly engaged, practiced, assessed, and explored.

The main advantage for cognitive science researchers in using robotic simulators is the possibility of reproducing the physics and dynamics of the robot and its interactions with the environment. It enables studying the behavior of different types of embodied agents without facing in advance the problem of building and maintaining a complex hardware device. Often, the simulator becomes a tool to test and validate an algorithm before porting it on a real robot. Mar, Tikhonoff, and Natale (2017) proposed a framework to identify the affordance properties of objects, with the goal of predicting the effect of the actions performed while using a novel tool. Their experiments were carried out on both the iCub robot and its simulator. The simulator was used to test in advance the effectiveness of the proposed framework and to automate extensive experiments on a large set of objects, which would have been tedious to perform on the real setup. When the case-study scenario makes real-world experimentations unfeasible or too expensive, the simulation becomes a valid alternative option. This is the case, for example, of social skills analyses, in which a proper experimental environment should include humans and be able to model dynamic interactions between humans, the robot, and the environment. Truong and Ngo (2017) solved the issue by simulating in Gazebo an office scenario, complete with doors, objects, and people. A socially aware mobile robot can navigate around the office with the goal of detecting humans, identifying their social state, and defining an approaching strategy.

Another advantage of simulators is that they allow many experiments to be conducted with the robot by varying its morphology and sensors without the need to develop these corresponding features in hardware. For example, Luo et al. (2018) proposed an infant-inspired framework for a robot to acquire reaching abilities. They used a simulation in Gazebo to evaluate the performance of the framework in dealing with diverse cases. They simulated two versions of the same robot with different arm lengths in order to imitate the

growth of the infants during the learning phase and the different uses of tools that can result.

Many studies in cognitive robotics are HRI studies. From this perspective, it is fundamental to include the human subject in the experimental environment. To do this, the simulation framework can be adapted to allow a user to interact with the simulation. For example, the user can send vocal inputs to the simulated robot as in Rossi, Staffa, and Rossi (2016), wherein the authors used V-REP to create a multirobot architecture, guided by a human operator, and analyzed how the vocal interaction evolved. Alternatively, a camera sensor can be used to monitor the movements of a human user and translate their gestures into commands for the robot. Caccavale and Finzi (2017) simulated both a Kuka omnirobot and a user in V-REP. The simulated user interacted with the robot by reproducing the gestures of a real human operator, whose movements were detected and recognized through an RGB-D sensor. A third way is to use the keyboard to generate events in the simulated environment, as in Pinto, Kuo, and Nikolaidis (2019). They used a reinforcement-learning framework in which a robotic arm collected data to learn a manipulation task while a human acted as an adversary in its learning process. The experiments were performed in Mujoco, where the user disturbed the interaction of the robot with the objects by applying force to the objects through the keyboard.

Since the 1980s, simulators have been part of the tools used for robotics research (Chan, Weston, and Case 1988). In the 2000s the interest in tools for robotics simulation and software development grew further, also thanks to the launch of tools such as Microsoft Robotics Developer Studio (Gates 2008; Jackson 2007) and USARSim (Carpin et al. 2007). The capability of simulators to reproduce the real world—both in terms of physics and photorealism—has been constantly improving. Recent progress obtained with large-scale training techniques, such as deep learning and reinforcement learning, have made simulation even more relevant. Deep architectures need to be trained on massive amounts of data in order to learn an effective and generalized representation of the world or effective control policies. The possibility of generating data with a simulated environment allows for faster data acquisition, without the need for a human operator to supervise the procedure and while avoiding damages to the real setup. This has been shown to work well in perception using a mix of data augmentation and photorealistic rendering to solve the problem of six-dimensional object pose estimation (Tremblay et al. 2018). The model developed by Mahler et al. (2017) can perform accurate precision grasps of many different objects by being trained on millions of depth images and grasp poses generated in simulation. Nowadays, a lot of effort is spent on the use of simulators to transfer skills and abilities learned in the simulated environment to the real-world system. This is particularly useful for reinforcement learning–based approaches in which the learning process may require months of real-world interaction, with the risk of damaging the robot, whereas in simulation it can be speeded up using modern parallel computing. However, the effectiveness of this approach is not straightforward due to the so-called reality gap, the discrepancy between reality and simulation that prevents simulated robotic experience from directly enabling effective real-world performance. A possible solution to this problem is to execute multiple simulations in which some of the parameters are randomized so the system can learn more robust control policies. An example is the recent work of Open AI, in which

a robot learns in-hand object manipulation with a dexterous hand (Andrychowicz et al. 2020).

In the context of cognitive robotics, the simulated task must support the underlying psychological and cognitive operations employed in performing the real-world task to ensure that the transfer effectively occurs. Many recent works address the problem, proposing promising ways to close the sim-to-real gap (Peng et al. 2018; Chebotar et al. 2019).

7.4.1 Types of Simulator

Simulators focus on different aspects of a robotic system. For example, it is possible to simulate how different parts of a robot deform given the forces that the robot exchanges with the environment. Another example is the simulation of physical quantities inside the robot, such as the temperature, or the current and voltage in the motors or boards. For researchers in cognitive robotics, the main focus is on real-time tools that can simulate full robot arms or humanoid robots in approximate real time on regular computers. In this context, “real time” means that one second of simulation takes approximately one second to be simulated, as opposed to specialized simulations that can be several orders of magnitude slower than real time. To run in real time, simulators typically disregard the simulation of fine details such as mechanical deformation or thermal propagation.

One of the major simplifications to achieve real time is to use multibody dynamics or rigid-body dynamics (Horak and Trinkle 2019; Featherstone and Orin 2016). This follows the assumption that robots are an assembly of multiple, perfectly rigid bodies, called links, interconnected by joints. Another simplifying assumption is to ignore the complex details of the actuators of the robot, whether electrical, hydraulic, or pneumatic, and just assume that it is possible to directly control the torque or force that the motors apply to the joints of the robot (Neunert, Boaventura, and Buchli 2016).

Available simulators can be classified into two main families: physics engines and simulation environments. Physics engines provide all the functionality necessary to simulate the physics of a system modeled as a rigid body, taking into account external forces and collisions with other simulated bodies. Simulation environments, instead, provide many other functionalities such as an integrated GUI, and they expose a user-friendly interface to one or more physics engines.

Several commercial and open-source physics engines are commonly used for robotics simulations. They are available as libraries for a given programming language, which is typically C/C++ given that performance is important in robot simulation. Examples of open-source physics engines are ODE, Bullet, and DART (these are discussed in detail in section 7.4.2).

In some cases, researchers use physics engine libraries directly to build their own simulators (an example is the iCub simulator based on the ODE physics engine (Tikhonoff et al. 2008), typically combining a physics engine with a rendering engine to visualize a three-dimensional model of the robot and the environment during the simulation. Simulation environments, on the other hand, are ready-to-use programs that permit the use of a physics engine, a rendering engine, and a user interface without the need to write code specific to each simulation scenario. In contrast, they provide description languages that allow the specification of the robot structure and the environment to be loaded through a file description. Examples of such simulation

environments are Gazebo, CoppeliaSim (formerly V-REP), Webots, and SIGVerse. These environments also support the ability to load code specific to given experiments in the form of custom plug-in systems or provide support for exposing the functionalities of the simulated robots using middleware interfaces or APIs, such as ROS/ROS2 or YARP (as an in-depth discussion of robot middleware like ROS is out of the scope of this chapter, we refer the reader to Kortenkamp, Simmons and Brugali [2016] and Magyar, Krizsán, and Sinčák [2015]).

For both physics engines and simulation environments, it is worth distinguishing two different use cases. In the first case, the user starts the simulation manually, as a real robot would be started, and then its execution continues in real time. In the second case, the user automatically runs multiple simulations at the same time, or multiple simulations over a long time—for example, for training a learning algorithm. These use cases respond to different needs of the users, and one environment can be optimized to provide more facilities for one use case or the other.

Another important aspect is the API exposed by the simulators to control the robot. In some cases the API is designed to replicate the interface of the real robot. This avoids the need to rewrite the control software when switching from the simulator to the real robot (proving in this way to be a digital twin of the real robot). This approach is followed, for example, by the iCub humanoid robot simulators (Tikhanoff et al. 2008; Hoffman et al. 2014).

Figure 7.3 shows examples of simulators from those built directly using the functionalities of a physics engine library to mature simulation environments able to reproduce realistic scenes with photorealistic rendering. In the next section, we discuss some of the physics engines and simulation environments we found in our survey and some we considered important given the current trend in robotics as of early 2020.

7.4.2 Available Simulators

Common open-source physics engines used extensively by the robotic community are Open Dynamic Engine (ODE; Smith 2020), Bullet (2020), and DART (2020). Excluding DART, these tools were originally developed for computer games and then adapted to work with robots. Nowadays these two domains have almost converged, providing at the same time accurate physics simulations and photorealistic rendering. A popular closed-source physics engine used extensively by the robot-learning community is MuJoCo (Roboti LLC 2020). For performance reasons, the physics engines are developed in low-level languages like C and C++, although they often provide bindings to other languages such as Python. The majority of both open-source and commercial simulation environments interface with at least one of the physics engines reported above.

Initial attempts to build robot simulators relied directly on the functionalities offered by a physics engine library. A notable example in this respect is the iCub ODE simulator (Tikhanoff et al. 2008), which used the ODE API to build a full simulation of the iCub robot, including all joints, the inertial sensors in the head, the cameras, and the facial expressions. The iCub ODE simulator also provides a software interface for position, velocity, and torque control. It also allows the loading of physical objects, directly from a configuration file or another software module using the YARP middleware. The iCub ODE simulator has been used in experiments with sensorimotor learning (Tommasino




Figure 7.3

Examples of simulators: (a) the iCub ODE simulator; (b) an example of simulation using Gazebo; and (c) the Isaac Sim. *Source:* Tikhonoff et al. 2008; (Hoffman et al. 2014; NVIDIA Corporation 2020). These examples show the evolution of simulation environments, from custom simulators programmed using the functionalities of a physics engine library to mature environments that allow the loading of complex scenes from description files and 3D models of objects, with high-fidelity rendering.

et al. 2019), tool affordances (Mar, Tikhonoff, and Natale 2017), estimation of affective states during face-to-face interaction (Boccignone et al. 2018), the study of computational models of development of language (Štěpánová et al. 2018), coordination of cognitive skills (Hwang and Tani 2018), and altruistic behavior (Baraglia, Nagai, and Asada 2016).

The main limitation of such simulators is their maintainability. Changes in the robot have to be propagated by modifying the simulation code; for this reason, it becomes difficult to support multiple robots or give the user the option to add new robots or objects to the simulation environment. Other benefits of using simulated environments, beyond improved user experience, include the possibility of extending the simulation with custom features and the opportunity to directly interact with simulated bodies from a graphic interface. In cognitive robotics research, particularly, interaction is of paramount importance. For these reasons, more recently, the community has shifted toward the adoption of simulation environments.

One of the most complete and enduring simulation environments is Gazebo (Koenig and Howard 2004), currently developed by the Open Software Robotics Foundation (2014)

and distributed as open-source software. The developers of Gazebo also proposed the SDF (Open Source Robotics Foundation 2019), an XML format, which describes robot models, the objects, and the environment in which the robot is deployed. With time the SDF was extended to describe all aspects that characterize a robot, static as well as dynamic objects, terrain, and lighting. Gazebo supports all the common sensors typically mounted on robots and allows developing software plug-ins to extend its capabilities. It also supports multiple physics engines such as ODE, Bullet, Simbody, and DART. Gazebo was initially developed for Linux, and more recently, it was extended to support Windows as well. The main advantage of Gazebo is its maturity and large community: Gazebo has been extensively used during the DARPA robotics challenge (Defense Advanced Research Projects Agency 2013) to simulate the ATLAS humanoid robot, and it is currently the simulator of choice for experiments on whole-body control and locomotion with the iCub robot (Hoffman et al. 2014). It has also been integrated as part of the Neurobotic Platform within the Human Brain Project (2018) to study models of the brain in simulated closed-loop systems (see, for example, Chen et al. 2019). Other examples are the simulation of a mobile platform for HRI studies (Truong and Ngo 2017) and the simulation of infant-like humanoid robots to investigate a developmental approach to learning reaching tasks (Luo et al. 2018).

CoppeliaSim (Rohmer, Singh, and Freese 2013), formerly V-REP, is a framework that, similarly to Gazebo, supports multiple physics engines, including Bullet and ODE. It is multiplatform, as it is distributed for macOS, Linux, and Windows. It can be used free of charge, but only in its educational version.

In our survey we found CoppeliaSim/V-REP used to simulate a custom robot system to study reinforcement learning for a domestic task (cleaning the table; Cruz et al. 2016), a multirobot mobile architecture (Rossi, Staffa, and Rossi 2016), and a Kuka omnibot in HRI settings (Caccavale and Finzi 2017).

Webots (Cyberbotics Ltd. 2020) is another simulation environment, which was developed in 1998 by the Swiss Federal Institute of Technology (EPFL) and became a commercial product of the EPFL spin-off company Cyberbotics (Olivier 2004). It was initially distributed as a closed-source application, and probably for this reason, its adoption suffered, especially when open-source alternatives like Gazebo gained popularity. With the release of the R2019a, Webots is being distributed with an open-source license. Webots is based on ODE; it provides a graphic interface that simplifies the design of the environment and achieves fast prototyping of robot systems starting from a set of sensors and actuators. It is multiplatform and runs on Windows, Linux, and macOS. In Pierris and Dahl (2017), Webots is employed to simulate a salamander-like robot to study an architecture for deriving novel skills by extending existing skills learned by demonstration.

Choregraphe (Pot et al. 2009) is the built-in programming application for Aldebaran robots, including NAO and Pepper. It allows the robot programmer to create animations, behaviors, and dialogues. Besides these programming capabilities, it lets users test the programmed behavior on a simulated robot, and for this reason, researchers who use Aldebaran robots often use it.

The MATLAB environment also provides a toolbox for robotic control and simulation, called the MATLAB Robotics System Toolbox (MathWorks 2020), which includes simulation tools integrated with Simulink and Simscape Mechanics. Furthermore, the MATLAB

Robotics System Toolbox offers out-of-the-box integration with the Gazebo simulator, permitting users to control Gazebo models with MATLAB and Simulink.

NVIDIA Corporation (2020) recently released Isaac Sim. It is part of the Isaac SDK, which additionally provides machine-learning algorithms and algorithms for motion planning, SLAM, and perception. Given the know-how of the company, it focuses on exploiting GPU acceleration for machine learning and simulation. For this reason and thanks to the support by NVIDIA, it is expected that Isaac Sim will soon be adopted by a large community.

A second categorization we discuss is between interactive or batch simulations. Interactive simulations are executed in real time, and robots belonging to the simulation are analogous to virtual replicas of real robots. Batch simulations, instead, are multiple instances of independent simulations that can run even faster than real time. The need for batch simulations has been driven by recent techniques proposed in reinforcement learning and their integration with deep learning, which demand training data sets obtained by running hundreds of thousands or even millions of actions. Key research areas in cognitive robotics such as imitation, behavior transfer, and knowledge acquisition could benefit from the usage of batch simulations. Simulation frameworks that support batch execution are Mujoco, PyBullet, CoppeliaSim, and Isaac Sim.

7.5 Conclusion

In this chapter we have provided an overview of the hardware platforms and simulation environments in use within the cognitive robotics community. We analyzed the literature in the field to identify the platforms commonly used in the community and described those found to be used most frequently.

Our main observation is that the community is widely fragmented. In our analysis of 337 papers, we found references to sixty-seven different hardware platforms and thirteen software simulators. Most of the platforms and simulators were used only once. Only a few platforms—namely, iCub, NAO, and Pepper—appeared to be adopted by a community of researchers. This is clearly a substantial problem because such a fragmentation in the community poses strong challenges in terms of experimental reproducibility and the sharing of code and research results. Yet the NAO and Pepper robots have been quite successful in building a community of researchers. We can speculate that this is due to their affordable cost and the fact that—being commercial products—they are more reliable than research platforms. Unfortunately, only a subset of the community—mostly researchers involved in HRI studies—has adopted these robots. The iCub robot, on the other hand, seems to support research that is more heterogeneous, ranging from HRI to sensorimotor learning, whole-body control, and learning. Yet the versatility of the platform comes with higher cost and complexity, which may reduce adoption, especially in groups that do not have a core expertise in robotics.

If we look at the field of robot grasping, however, we notice a different trend, in which platforms such as the Panda arm and the Baxter robot are becoming *de facto* standards for research. We believe this is because these platforms strike a good balance between cost, reliability, complexity, and the type of research questions they allow users to address. In this respect, it seems challenging for a single platform to serve the whole cognitive robotics

community. In any case it seems clear that the community would greatly benefit from a larger adoption of shared platforms and that new platforms able to meet at least a subset of the requirements of the researchers would have a large impact on the community.

In our survey we also noticed a mild interest in simulations, outlined by a very scattered adoption of simulation tools. We argued that a possible reason could be that simulators do not yet provide complex models of the environment and do not allow the modeling of realistic interactions with humans. In fact, since HRI is often bidirectional, a simulator for HRI should provide interfaces for the robot to receive input from the human and for the human to receive feedback (visual and acoustic, but also haptic) from the simulated robot. Advances in virtual and augmented reality technology may progressively fill the gap; however, their integration with robotic simulators has not been extensively explored yet.

Several novel applications of simulation tools could find applications in cognitive robotics, even if their use is not currently widespread. Recent technologies developed for *virtual reality* (VR), such as VR headsets, would allow human users to interact naturally with simulated robots, as is done, for example, in SIGVerse (Mizuchi and Inamura 2017). At the same time, various technologies that sense humans are maturing, and they could be used to reproduce the movement of the human inside the simulation. Examples of such technologies are body-tracking systems that rely on vision (e.g., CMU open pose; Cao et al. 2017) and sensorized suits that integrate information from a distributed network of inertial units and torque and pressure sensors (Latella et al. 2019).

Finally, for the subfield of physical HRI, a useful future application of simulation tools may be physical interfaces able to provide users with force feedback from the simulation, using haptic feedback devices (Hannaford and Okamura 2016). These will simulate not only visual interaction between the user and the robot but also physical contact arising from the interaction.

Additional Reading and Resources

- A recent, complete handbook on humanoid robotics, with specific sections on robot platforms (part II) and simulators (part VIII): Ambarish G., and V. Prahlaad, eds. 2019. *Humanoid Robotics: A Reference*. Netherlands: Springer.
- IEEE robots—your guide to robotics: <https://robots.ieee.org>.
- ROS robot operating system: <https://www.ros.org>.
- Official iCub website with links to robot simulator and middleware: <https://icub.iit.it/>.
- IEEE education resources in robotics: <https://www.ieee-ras.org/educational-resources-outreach/educational-material-in-robotics-and-automation>.

References

- Ahn, Jonggil, and Gerard Jounghyun Kim. 2018. “SPRinT: A Mixed Approach to a Hand-Held Robot Interface for Telepresence.” *International Journal of Social Robotics* 10 (4): 537–552.
- Andrychowicz, Marcin, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron et al. 2020. “Learning Dexterous In-Hand Manipulation.” *International Journal of Robotics Research* 39 (1): 3–20.
- Anzalone, Salvatore Maria, Giovanna Varni, Serena Ivaldi, and Mohamed Chetouani. 2017. “Automated Prediction of Extraversion during Human–Humanoid Interaction.” *International Journal of Social Robotics* 9 (3): 385–399.

- ATR-Creative. 2010. "Main Specifications of Robovie-R Ver.3." Accessed January 16, 2020. <http://atr-c.jp/robot/r3/robo-r3spec.html>.
- Autodesk Inc. 2020. "3DS Max." Accessed June 4, 2020. <https://www.autodesk.it/products/3ds-max/overview>.
- Baraglia, Jimmy, Yukie Nagai, and Minoru Asada. 2016. "Emergence of Altruistic Behavior through the Minimization of Prediction Error." *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 141–151.
- Benli, Emrah, Yuichi Motai, and John Rogers. 2017. "Human Behavior-Based Target Tracking with an Omnidirectional Thermal Camera." *IEEE Transactions on Cognitive and Developmental Systems* 11 (1): 36–50.
- Blender Foundation. 2020. "About." <https://www.blender.org/>.
- Boccignone, Giuseppe, Donatello Conte, Vittorio Cuculo, Alessandro D'Amelio, Giuliano Grossi, and Raffaella Lanzarotti. 2018. "Deep Construction of an Affective Latent Space via Multimodal Enactment." *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 865–880.
- Bullet (website). 2020. "Bullet Real-Time Physics Simulation." Accessed February 13, 2020. <https://pybullet.org>.
- Busch, Baptiste, Jonathan Grizou, Manuel Lopes, and Freck Stulp. 2017. "Learning Legible Motion from Human-Robot Interactions." *International Journal of Social Robotics* 9 (5): 765–779.
- Caccavale, Riccardo, and Alberto Finzi. 2017. "Flexible Task Execution and Attentional Regulations in Human-Robot Interaction." *IEEE Transactions on Cognitive and Developmental Systems* 9 (1): 68–79.
- Cao, Hoang-Long, Greet Ven de Perre, James Kennedy, Emmanuel Senft, Pablo Gómez Esteban, Alberto De Beir, Ramona Simut, Tony Belpaeme, Dirk Lefeber, and Bram Vanderborght. 2019. "A Personalized and Platform-Independent Behavior Control System for Social Robots in Therapy: Development and Applications." *IEEE Transactions on Cognitive and Developmental Systems* 11 (3): 334–346.
- Cao, Zhe, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. "Realtime Multi-person 2D Pose Estimation Using Part Affinity Fields." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7291–7299. New York: IEEE.
- Caré, Xavier. 2015. "Innorobo 2015—Kuka Robotics." July 2. Accessed May 21, 2020. https://commons.wikimedia.org/wiki/File:Innorobo_2015_-_Kuka_Robotics.JPG.
- Carpin, Stefano, Mike Lewis, Jijun Wang, Stephen Balakirsky, and Chris Scrapper. 2007. "USARSim: A Robot Simulator for Research and Education." In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, 1400–1405. New York: IEEE.
- Çelikkanat, Hande, Güner Orhan, Nicolas Pugeault, Frank Guerin, Erol Şahin, and Sinan Kalkan. 2016. "Learning Context on a Humanoid Robot Using Incremental Latent Dirichlet Allocation." *IEEE Transactions on Cognitive and Developmental Systems* 8 (1): 42–59.
- Chan, S. F. R. H Weston, and K. Case. 1988. "Robot Simulation and Off-Line Programming." *Computer-Aided Engineering Journal* 5 (4): 157–162.
- Chandra, Shruti, Pierre Dillenbourg, and Ana Paiva. 2020. "Children Teach Handwriting to a Social Robot with Different Learning Competencies." *International Journal of Social Robotics* 12 (3):721–748.
- Chebotar, Yevgen, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. 2019. "Closing the Sim-to-Real Loop: Adapting Simulation Randomization with Real World Experience." In *2019 International Conference on Robotics and Automation*, 8973–8979. New York: IEEE.
- Chen, Guang, Zhenshan Bing, Florian Röhrbein, Jörg Conradt, Kai Huang, Long Cheng, Zhuangyi Jiang, and Alois Knoll. 2019. "Toward Brain-Inspired Learning with the Neuromorphic Snake-Like Robot and the Neuro-robotic Platform." *IEEE Transactions on Cognitive and Developmental Systems* 11 (1): 1–12.
- Claret, Josep-Arnau, Gentiane Venture, and Luis Basañez. 2017. "Exploiting the Robot Kinematic Redundancy for Emotion Conveyance to Humans as a Lower Priority Task." *International journal of social robotics* 9 (2): 277–292.
- Coppelia Robotics GmbH. 2020. "CoppeliaSim." <https://www.coppeliarobotics.com/>.
- Cruz, Francisco, Sven Magg, Cornelius Weber, and Stefan Wermter. 2016. "Training Agents with Interactive Reinforcement Learning and Contextual Affordances." *IEEE Transactions on Cognitive and Developmental Systems* 8 (4): 271–284.
- Cyberbotics Ltd. 2020. "Webots Open Source Robot Simulator." Accessed February 13, 2020. <https://cyberbotics.com/>.
- DART (website). 2020. "Dynamic Animation and Robotics Toolkit." Accessed February 13, 2020. <https://dartsim.github.io/>.
- David, Daniel O., Cristina A. Costescu, Silviu Matu, Aurora Szentagotai, and Anca Dobrea. 2018. "Developing Joint Attention for Children with Autism in Robot-Enhanced Therapy." *International Journal of Social Robotics* 10 (5): 595–605.
- Defense Advanced Research Projects Agency. 2013. "DARPA Robotics Challenge." <https://www.darpa.mil/program/darpa-robotics-challenge>.

- Donner, Philine, Franz Christange, Jing Lu, and Martin Buss. 2017. “Cooperative Dynamic Manipulation of Unknown Flexible Objects.” *International Journal of Social Robotics* 9 (4): 575–599.
- Energy.gov. 2013. “AEMC Summit.” December 23. Accessed May 21, 2020. [https://commons.wikimedia.org/wiki/File:AEMC_Summit_\(11343380734\).jpg](https://commons.wikimedia.org/wiki/File:AEMC_Summit_(11343380734).jpg).
- Farsa Sourceforge (website). 2020. “Framework for Autonomous Robotics Simulation and Analysis.” Accessed June 4, 2020. <https://sourceforge.net/projects/farsa/>.
- Featherstone, Roy, and David E. Orin. 2016. “Dynamics.” In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Kathib Oussama, 37–66. Cham, Switzerland: Springer.
- Fitzgerald, Cliff. 2013. “Developing Baxter.” In *2013 IEEE Conference on Technologies for Practical Robot Applications*, 1–6. New York: IEEE.
- Fondazione Istituto Italiano di Tecnologia. n.d. <https://github.com/robotology/idyntree>. Accessed April 1, 2018.
- Förster, Frank, Joe Saunders, and Chrystopher L. Nehaniv. 2018. “Robots That Say ‘No’ Affective Symbol Grounding and the Case of Intent Interpretations.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 530–544.
- Franka Emika GmbH. 2020. “Introducing the Franka Emika Robot.” Accessed February 5, 2020. <https://www.franka.de/>.
- Gates, Bill. 2008. “A Robot in Every Home.” February 1. Accessed June 4, 2020. <https://www.scientificamerican.com/article/a-robot-in-every-home-2008-02/>.
- Generation Robots. 2020. “PANDA Robotic Arm.” Accessed February 13, 2020. <https://www.generationrobots.com/en/403317-panda-robotic-arm.html>.
- Giagkos, Alexandros, Daniel Lewkowicz, Patricia Shaw, Suresh Kumar, Lee Mark, and Qiang Shen. 2017. “Perception of Localized Features during Robotic Sensorimotor Development.” *IEEE Transactions on Cognitive and Developmental Systems* 9 (2): 127–140.
- Glover, Arren J., and Gordon F. Wyeth. 2018. “Toward Lifelong Affordance Learning Using a Distributed Markov Model.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (1): 44–55.
- Gouaillier, David, Vincent Hugel, Pierre Blazeovic, Chris Kilner, Jérôme Monceaux, Pascal Lafourcade, Brice Marnier, Julien Serre, and Bruno Maisonnier. 2009. “Mechatronic Design of NAO Humanoid.” In *2009 IEEE International Conference on Robotics and Automation*, 769–774. New York: IEEE.
- Hannaford, Blake, and Allison M. Okamura. 2016. “Haptics.” In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Kathib Oussama, 1063–1084. Cham, Switzerland: Springer.
- Herath, Damith C., Elizabeth Jochum, and Evgenios Vlachos. 2018. “An Experimental Study of Embodied Interaction and Human Perception of Social Presence for Interactive Robots in Public Settings.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 1096–1105.
- Hirano, Takahiro, Masahiro Shiomi, Takamasa Iio, Mitsuhiro Kimoto, Ivan Tanev, Katsunori Shimohara, and Norihiro Hagita. 2018. “How Do Communication Cues Change Impressions of Human-Robot Touch Interaction?” *International Journal of Social Robotics* 10 (1): 21–31.
- Hirzinger, Gerd, Norbert Sporer, Alin Albu-Schaffer, M. Hahnle, Rainer Krenn, Antonio Pascucci, and Markus Schedl. 2002. “DLR’s Torque-Controlled Light Weight Robot III—Are We Reaching the Technological Limits Now?” In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*. Cat. No. 02CH37292. Vol. 2, 1710–1716. New York: IEEE.
- Hoffman, Enrico Mingo, Silvio Traversaro, Alessio Rocchi, Mirko Ferrati, Alessandro Settini, Francesco Romano, Lorenzo Natale, Antonio Bicchi, Francesco Nori, and Nikos G. Tsagarakis. 2014. “Yarp Based Plugins for Gazebo Simulator.” In *International Workshop on Modelling and Simulation for Autonomous Systems*, 333–346. Cham, Switzerland: Springer.
- Hoffmann, Matej, Zdeněk Straka, Igor Farkaš, Michal Vavrečka, and Giorgio Metta. 2018. “Robotic Homunculus: Learning of Artificial Skin Representation in a Humanoid Robot Motivated by Primary Somatosensory Cortex.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 163–176.
- Horak, Peter C., and Jeff C. Trinkle. 2019. “On the Similarities and Differences among Contact Models in Robot Simulation.” *IEEE Robotics and Automation Letters* 4 (2): 493–499.
- Human Brain Project. 2018. *HBP Neurorobotics Platform*. Accessed February 13, 2020. <https://www.neurorobotics.net/>.
- Hwang, Jungsik, and Jun Tani. 2018. “Seamless Integration and Coordination of Cognitive Skills in Humanoid Robots: A Deep Learning Approach.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 345–358.
- IEEE Spectrum*. 2020a. “Franka: A Robot Arm That’s Safe, Low Cost, and Can Replicate Itself.” Accessed February 13, 2020. <https://spectrum.ieee.org/robotics/industrial-robots/franka-a-robot-arm-thats-safe-low-cost-and-can-replicate-itself>.

- IEEE Spectrum*. 2020b. “Robots: Your Guide to the World of Robotics.” Accessed May 5, 2020. <https://robots.ieee.org/robots/pioneer/>.
- Iio, Takamasa, Masahiro Shiomi, Kazuhiko Shinozawa, Katsunori Shimohara, Mitsunori Miki, and Norihiro Hagita. 2015. “Lexical Entrainment in Human Robot Interaction.” *International Journal of Social Robotics* 7 (2): 253–263.
- Ims. 2017. “Franka Emika.” March 23. Accessed May 21, 2020. <https://commons.wikimedia.org/w/index.php?curid=57761214>.
- Ivaldi, Serena, Sebastien Lefort, Jan Peters, Mohamed Chetouani, Joelle Provasi, and Elisabetta Zibetti. 2017. “Towards Engagement Models That Consider Individual Factors in HRI: On the Relation of Extroversion and Negative Attitude towards Robots to Gaze and Speech during a Human-Robot Assembly Task.” *International Journal of Social Robotics* 9 (1): 63–86.
- Izui, Takamune, and Gentiane Venture. 2020. “Correlation Analysis for Predictive Models of Robot User’s Impression: A Study on Visual Medium and Mechanical Noise.” *International Journal of Social Robotics* 12 (2): 425–439.
- Jackson, Jared. 2007. “Microsoft Robotics Studio: A Technical Introduction.” *IEEE Robotics and Automation Magazine* 14 (4): 82–87.
- Jones, Aidan, and Ginevra Castellano. 2018. “Adaptive Robotic Tutors That Support Self-Regulated Learning: A Longer-Term Investigation with Primary School Children.” *International Journal of Social Robotics* 10 (3): 357–370.
- Karunarathne, Deneth, Yoichi Morales, Takayuki Kanda, and Hiroshi Ishiguro. 2018. “Model of Side-by-Side Walking without the Robot Knowing the Goal.” *International Journal of Social Robotics* 10 (4): 401–420.
- Katz, Garrett, Di-Wei Huang, Theresa Hauge, Rodolphe Gentili, and James Reggia. 2018. “A Novel Parsimonious Cause-Effect Reasoning Algorithm for Robot Imitation and Plan Recognition.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 177–193.
- Khamassi, Mehdi, George Velentzas, Tsitsimis Theodore, and Costas S. Tzafestas. 2019. “Robot Fast Adaptation to Changes in Human Engagement during Simulated Dynamic Social Interaction with Active Exploration in Parameterized Reinforcement Learning.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 881–893.
- Koenig, Nathan, and Andrew Howard. 2004. “Design and Use Paradigms for Gazebo, an Open-Source Multi-robot Simulator.” In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Cat. No. 04CH37566. Vol. 3, 2149–2154. New York: IEEE.
- Kortenkamp, David, Reid Simmons, and Davide Brugali. 2016. “Robotic Systems Architectures and Programming.” In *Springer Handbook of Robotics*, edited by Siciliano Bruno and Khatib Ossama, 283–306. Cham, Switzerland: Springer.
- Kwak, Sonya S., Jun San Kim, and Jung Ju Choi. 2017. “The Effects of Organism versus Object-Based Robot Design Approaches on the Consumer Acceptance of Domestic Robots.” *International Journal of Social Robotics* 9 (3): 359–377.
- Latella, Claudia, Silvio Traversaro, Diego Ferigo, Yeshasvi Tirupachuri, Lorenzo Rapetti, Francisco Javier Andrade Chavez, Francesco Nori, and Daniele Pucci. 2019. “Simultaneous Floating-Base Estimation of Human Kinematics and Joint Torques.” *Sensors* 19 (12).
- Laurettil, Clemente, Francesca Cordella, and Loredana Zollo. 2019. “A Hybrid Joint/Cartesian DMP-Based Approach for Obstacle Avoidance of Anthropomorphic Assistive Robots.” *International Journal of Social Robotics* 11 (5): 783–796.
- Liu, Phoebe, Dylan F. Glas, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2017. “A Model for Generating Socially-Appropriate Deictic Behaviors towards People.” *International Journal of Social Robotics* 9 (1): 33–49.
- Liu, Rui, and Xiaoli Zhang. 2016. “Understanding Human Behaviors with an Object Functional Role Perspective for Robotics.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (2): 115–127.
- Luo, Dingsheng, Fan Hu, Tao Zhang, Yian Deng, and Xihong Wu. 2018. “How Does a Robot Develop Its Reaching Ability Like Human Infants Do?” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 795–809.
- Magyar, Gergely, Peter Krizsán, and Zoltán Sinčák. 2015. “Comparison Study of Robotic Middleware for Robotic Applications.” In *Emergent Trends in Robotics and Intelligent Systems*, edited by P. Hartono, M. Virčíková, J. Vaščák, and R. Jakša, 121–128. Cham, Switzerland: Springer.
- Mahler, Jeffrey, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, and Ken Goldberg. 2017. “Dex-net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics.” ArXiv preprint: 1703.09312.
- Mar, Tanis, Vadim Tikhonoff, and Lorenzo Natale. 2017. “What Can I Do with This Tool? Self-Supervised Learning of Tool Affordances from Their 3-D Geometry.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 595–610.

- MathWorks. 2020. “Robotics System Toolbox.” Accessed May 26, 2020. <https://www.mathworks.com/products/robotics.html>.
- Meola, Valentina Cristina, Daniele Caligiore, Valerio Sperati, Loredana Zollo, Anna Lisa Ciancio, Fabrizio Taffoni, Eugenio Guglielmelli, and Gianluca Baldassarre. 2016. “Interplay of Rhythmic and Discrete Manipulation Movements during Development: A Policy-Search Reinforcement-Learning Robot Model.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 152–170.
- Mizuchi, Yoshiaki, and Tetsunari Inamura. 2017. “Cloud-Based Multimodal Human-Robot Interaction Simulator Utilizing ROS and Unity Frameworks.” In *IEEE/SICE International Symposium on System Integration*. New York: IEEE.
- Moulin-Frier, Clément, Tobias Fischer, Maxime Petit, Grégoire Pointeau, Jordi-Ysard Puigbo, Ugo Pattacini, Sock Ching Low et al. 2018. “DAC-h3: A Proactive Robot Cognitive Architecture to Acquire and Express Knowledge about the World and the Self.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 1005–1022.
- Murata, Shingo, Yuxi Li, Hiroaki Arie, Tetsuya Ogata, and Shigeki Sugano. 2018. “Learning to Achieve Different Levels of Adaptability for Human-Robot Collaboration Utilizing a Neuro-dynamical System.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 712–725.
- Natale, Lorenzo, Chiara Bartolozzi, Francesco Nori, Giulio Sandini, and Giorgio Metta. 2017. “iCub.” In *Humanoid Robotics: A Reference*, edited by P. Vadakkepat, A. Goswami, and P. Vadakkepat, 1–33. Dordrecht: Springer.
- Natale, Lorenzo, Ali Paikan, Marco Randazzo, and Daniele E. Domenichelli. 2016. “The iCub Software Architecture: Evolution and Lessons Learned.” *Frontiers in Robotics and AI* 3 (25).
- Neunert, Michael, Thiago Boaventura, and Jonas Buchli. 2016. “Why Off-the-Shelf Physics Simulators Fail in Evaluating Feedback Controller Performance—a Case Study for Quadrupedal Robots.” In *Advances in Cooperative Robotics: Proceedings of the 19th International Conference on CLAWAR 2016*, edited by M. O. Tokhi and G. S. Virk, 464–472. Hackensack, NJ: World Scientific.
- Nomura, Tatsuya, and Takayuki Kanda. 2015. “Influences of Evaluation and Gaze from a Robot and Humans’ Fear of Negative Evaluation on Their Preferences of the Robot.” *International Journal of Social Robotics* 7 (2): 155–164.
- NVIDIA Corporation. 2020. “NVIDIA Isaac Sim. Accessed February 13, 2020. <https://developer.nvidia.com/isaac-sim>.
- Olivier, Michael. 2004. “WebotsTM: Professional Mobile Robot Simulation.” *International Journal of Advanced Robotic Systems* 1 (1): 40–43.
- OpenSim (website). 2020. Accessed June 4, 2020. <http://simtk.org/projects/opensim>.
- Open Software Robotics Foundation. 2014. “GAZEBO Robot Simulation Made Easy.” Accessed February 13, 2020. <http://gazebosim.org/>.
- Open Source Robotics Foundation. 2019. “SDF Describe Your World.” Accessed February 13, 2020. <http://sdformat.org/>.
- Pandey, Amit Kumar, and Rodolphe Gelin. 2018. “A Mass-Produced Sociable Humanoid Robot: Pepper: The First Machine of Its Kind.” *IEEE Robotics Automation Magazine* 25 (3): 40–48.
- Park, Jun-Cheol, Dae-Shik Kim, and Yukie Nagai. 2017. “Learning for Goal-Directed Actions Using RNNPB: Developmental Change of ‘What to Imitate.’” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 545–556.
- Parmiggiani, Alberto, Marco Maggiali, Lorenzo Natale, Francesco Nori, Alexander Schmitz, Nikos Tsarakakis, Jose Santos Victor, Francesco Becchi, Giulio Sandini, and Giorgio Metta. 2012. “The Design of the iCub Humanoid Robot.” *International Journal of Humanoid Robotics* 9 (4): 1250027.
- Peng, Xue Bin, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. 2018. “Sim-to-Real Transfer of Robotic Control with Dynamics Randomization.” In *2018 IEEE International Conference on Robotics and Automation*, 1–8. New York: IEEE.
- Petit, Maxime, Tobias Fischer, and Yiannis Demiris. 2016. “Lifelong Augmentation of Multimodal Streaming Autobiographical Memories.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 201–213.
- Pierris, Georgios, and Torbjørn S. Dahl. 2017. “Learning Robot Control Using a Hierarchical SOM-Based Encoding.” *IEEE Transactions on Cognitive and Developmental Systems* 9 (1): 30–43.
- Pinto, Lerrel, C-C. Jay Kuo, and Stefanos Nikolaidis. 2019. “Robot Learning via Human Adversarial Games.” ArXiv preprint: 1903.00636.
- Pot, Emmanuel, Jérôme Monceaux, Rodolphe Gelin, and Bruno Maisonier. 2009. “Choregraphe: A Graphical Tool for Humanoid Robot Programming.” In *RO-MAN 2009—the 18th IEEE International Symposium on Robot and Human Interactive Communication*, 46–51. New York: IEEE.
- Ribes, Arturo, Jesus Cerquides, Yiannis Demiris, and Ramon Lopez de Mantaras. 2016. “Active Learning of Object and Body Models with Time Constraints on a Humanoid Robot.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (1): 26–41.

- RobotCub@Home. 2018. “SoftBank Pepper Conditions 2018.” Accessed May 5, 2020. http://www.robocupathome.org/athome-spl/pepper_conditions_18.
- Robotics Business Review*. 2015. “Is Sale of Universal Robots Classic Innovator’s Dilemma?” Accessed February 13, 2020. https://www.roboticsbusinessreview.com/manufacturing/is_sale_of_universal_robots_classic_innovators_dilemma/.
- Roboti LLC. 2020. “MuJoCo: Advanced Physics Simulation.” Accessed February 13, 2020. <http://www.mujooco.org/>.
- RobotShop Community. 2020. “Robovie R3.” Accessed February 13, 2020. <https://www.robotshop.com/community/blog/show/robovie-r3>.
- Rohmer, Eric, Surya P. N. Singh, and Marc Freese. 2013. “CoppeliaSim (formerly V-REP): A Versatile and Scalable Robot Simulation Framework.” In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. New York: IEEE.
- Rossi, Alessandra, Mariacarla Staffa, and Silvia Rossi. 2016. “Supervisory Control of Multiple Robots through Group Communication.” *IEEE Transactions on Cognitive and Developmental Systems* 9 (1): 56–67.
- Santucci, Vieri Giuliano, Gianluca Baldassarre, and Marco Mirolli. 2016. “GRAIL: A Goal-Discovering Robotic Architecture for Intrinsically-Motivated Learning.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 214–231.
- Saputra, Azhar Aulia, János Botzheim, and Naoyuki Kubota. 2019. “Evolving a Sensory-Motor Interconnection Structure for Adaptive Biped Robot Locomotion.” *IEEE Transactions on Cognitive and Developmental Systems* 11 (2): 244–256.
- SDU Robotics. 2020. “RobWork.” Accessed June 4, 2020. <https://www.robwork.org/>.
- SIGVerse (website). 2020. “SIGVerse.” Accessed June 4, 2020. <http://www.sigverse.org/wiki/en/>.
- Smashing Robotics. 2016. “Thirteen Advanced Humanoid Robots for Sale Today.” Accessed February 13, 2020. <https://www.smashingrobotics.com/thirteen-advanced-humanoid-robots-for-sale-today/>.
- SMILE (GitHub page). 2020. “SMILE: Simulator for Maryland Imitation Learning Environment.” Accessed June 4, 2020. <https://github.com/dwhuang/SMILE>.
- Smith, Russ. 2020. “Open Dynamics Engine.” Accessed February 13, 2020. <https://www.ode.org/>.
- Sony Corporation. 2020. “Aibos History.” Accessed February 13, 2020. <http://www.sony-aibo.com/aibos-history/>.
- Stage (GitHub) page. 2020. “The Stage Simulator.” Accessed June 4, 2020. <https://github.com/rtv/Stage>.
- Štěpánová, Karla, Frederico Belmonte Klein, Angelo Cangelosi, and Michal Vavrečka. 2018. “Mapping Language to Vision in a Real-World Robotic Scenario.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 784–794.
- Tan, Jie, Tingnan Zhang, Erwin Coumans, Atıl İscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. 2018. “Sim-to-Real: Learning Agile Locomotion for Quadruped Robots.” ArXiv preprint: 1804.10332.
- TechCrunch. 2015. “The Enterprise Model of Pepper, SoftBank’s Robot, Will Cost \$440 A Month to Rent.” July 30. Accessed May 22, 2020. <https://techcrunch.com/2015/07/30/pepper-earns-its-keep/>.
- Tikhonoff, Vadim, Angelo Cangelosi, Paul Fitzpatrick, Giorgio Metta, Lorenzo Natale, and Francesco Nori. 2008. “An Open-Source Simulator for Cognitive Robotics Research: The Prototype of the iCub Humanoid Robot Simulator.” In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, 57–61. New York: Association for Computing Machinery.
- Tokumeigakarinoashima. 2014. “SoftBank Pepper.” July 18. Accessed May 21, 2020. https://commons.wikimedia.org/wiki/File:SoftBank_pepper.JPG.
- Tommasino, Paolo, Daniele Caligiore, Marco Mirolli, and Gianluca Baldassarre. 2019. “A Reinforcement Learning Architecture That Transfers Knowledge between Skills When Solving Multiple Tasks.” *IEEE Transactions on Cognitive and Developmental Systems* 11 (2): 292–317.
- Tremblay, Jonathan, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. 2018. “Deep Object Pose Estimation for Semantic Robotic Grasping of Household Objects.” ArXiv preprint: 1809.10790.
- Truong, Xuan-Tung, and Trung-Dung Ngo. 2017. “To Approach Humans? A Unified Framework for Approaching Pose Prediction and Socially Aware Robot Navigation.” *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 557–572.
- Tsiourti, Christiana, Astrid Weiss, Katarzyna Wac, and Markus Vincze. 2019. “Multimodal Integration of Emotional Signals from Voice, Body, and Context: Effects of (In) Congruence on Emotion Recognition and Attitudes towards Robots.” *International Journal of Social Robotics* 11 (4): 555–573.
- Ugur, Emre, and Justus Piater. 2017. “Emergent Structuring of Interdependent Affordance Learning Tasks Using Intrinsic Motivation and Empirical Feature Selection.” *IEEE Transactions on Cognitive and Developmental Systems* 9 (4): 328–340.

- Wang, Jiru, Vui Ann Shim, Rui Yan, Huajin Tang, and Fuchun Sun. 2019. "Automatic Object Searching and Behavior Learning for Mobile Robots in Unstructured Environment by Deep Belief Networks." *IEEE Transactions on Cognitive and Developmental Systems* 11 (3): 395–404.
- Wang, Jiuguang. 2008. "ActivMedia Pioneer 3-AT Robot." May 8. Accessed May 21, 2020. https://commons.wikimedia.org/wiki/File:ActivMedia_Pioneer_3-AT_robot.jpg.
- Wieser, Erhard, and Gordon Cheng. 2018. "A Self-Verifying Cognitive Architecture for Robust Bootstrapping of Sensory-Motor Skills via Multipurpose Predictors." *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 1081–1095.
- Wikipedia. 2020. "Baxter (Robot)." Accessed February 13, 2020. [https://en.wikipedia.org/wiki/Baxter_\(robot\)](https://en.wikipedia.org/wiki/Baxter_(robot)).
- Willemse, Cesco, and Agnieszka Wykowska. 2019. "In Natural Interaction with Embodied Robots, We Prefer It When They Follow Our Gaze: A Gaze-Contingent Mobile Eyetracking Study." *Philosophical Transactions of the Royal Society B* 374 (1771): 20180036.
- Yang, Chenguang, Chuize Chen, Ning Wang, Zhaojie Ju, Jian Fu, and Min Wang. 2018. "Biologically Inspired Motion Modeling and Neural Control for Robot Learning from Demonstrations." *IEEE Transactions on Cognitive and Developmental Systems* 11 (2): 281–291.
- Zambelli, Martina, and Yiannis Demiris. 2017. "Online Multimodal Ensemble Learning Using Self-Learned Sensorimotor Representations." *IEEE Transactions on Cognitive and Developmental Systems* 9 (2): 113–126.
- Zlotowski, Jakub, Hidenobu Sumioka, Friederike Eyssele, Shuichi Nishio, Christoph Bartneck, and Hiroshi Ishiguro. 2018. "Model of Dual Anthropomorphism: The Relationship between the Media Equation Effect and Implicit Anthropomorphism." *International Journal of Social Robotics* 10 (5): 701–714.

8

Biomimetic Skin

Markellos Ntagios, Oliver Ozioko, and Ravinder Dahiya

8.1 Introduction

Replicating the fundamental characteristics of biological organs to develop their artificial equivalents and using them in robotic platforms is an area that is attracting significant interest through topics such as soft robotics; electronic skin, or eskin; and bionic limbs (Dahiya 2019; Dahiya, Akinwande et al. 2019; Dahiya, Yogeswaran et al. 2019; Soni and Dahiya 2020). The interest in this field is also fueled by the new and emerging applications of robots in areas such as smart factories and ambient assisted living, where safe and intelligent human-robot interaction is necessary. For robotic systems to move from industrial environments to home and urban or social areas, it is critical for them to have human-skin-like capabilities in order to enable safe human and robot interaction (Argall and Billard 2010; Dahiya et al. 2013). Robotic systems need to function close to humans for this to be achieved; therefore, the equivalents of human organs are needed for robots. Pacemakers and cochlear implants are some of the artificial organs developed in the past. The successful commercialization of some of the bionic organs such as electronic noses and ears and bionic eyes has encouraged researchers to explore more artificial organs—for example, eskin or tactile skin. This progress is also supported by technological advances in soft and flexible electronics (Gupta, Navaraj, et al. 2018; Núñez, Manjakkal, and Dahiya 2019), which could allow tactile skin to conform to curved surfaces (Hammock et al. 2013; Dahiya, Yogeswaran, et al. 2019); artificial muscles (Roche et al. 2014); and computation including artificial intelligence (AI; Decherchi et al. 2011; Luo et al. 2017; Navaraj et al. 2017). However, current advances still fall short of leading us to the functionalities offered by human skin. A deeper look at the sensory mechanisms in the human body shows the importance of the “sense of touch” in wide-ranging tasks such as the assessment of various properties of real-world objects and their handling. The size, shape, texture, temperature, surface roughness, hardness, softness, curvature, and more can all be assessed by touching. To determine such parameters, the human skin has different types of receptors that are distributed nonuniformly throughout the body, as discussed in the next section (Dahiya, Metta et al. 2009; Dahiya and Valle 2013; Dahiya, Mittendorfer et al. 2013; Yogeswaran, Dang et al. 2015). These receptors are embedded at different depths in the soft skin. It is challenging to realize an artificial skin with this level of complexity, especially when soft electronics technology is still at an early stage of development. Furthermore,

the sensing feature of skin is intimately connected with computation, actuation, and energy (Soni and Dahiya 2020). An eskin with tightly coupled sensing, actuation, computation, and energy devices over a large area will be hugely beneficial for robotics as well as other emerging areas such as autonomous vehicles, tactile internet (Simsek et al. 2016), and augmented/virtual reality in which intelligent interaction is desired. This chapter focuses on a new perspective related to eskin or tactile skin and presents some case studies. Section 8.1 presents a new approach for obtaining sensorized complex structures such as robotic or prosthetic hands. The advanced multimaterial, three-dimensional (3D) printing approach and the innovative designs used to realize the robotic hand with embedded sensors, actuators, and electronics are presented in section 8.3. Section 8.4 presents another case study in which different types of transducers (piezoelectric and capacitive) have been stacked to obtain the FA (fast-adapting) and SA (slow-adapting) receptors' equivalents of the human skin. The presented sensor stack is expected to allow eskin to detect both static and dynamic tactile or contact events. Furthermore, the machine-learning approach has been used to demonstrate the texture-detection capability of the presented sensor stack. Last, section 8.5 describes a new soft sensor device with a tightly coupled touch sensor and actuator. Altogether, these case studies show how eskin research is advancing toward tightly coupled sensing, actuation, and computation.

8.2 Tactile Sensing

The human skin is the largest organ of the human body. It comprises multiple mechanoreceptors, classified into two major categories (FAs and SAs) based on their response (table 8.1). The FA mechanoreceptors (Meissner's corpuscles and Pacinian corpuscles) are responsible for the detection of dynamic contact force/pressure applied to the human skin. They respond to slippage, to high-frequency vibration, and to the onset and offset of stimulation. On the other hand, SA mechanoreceptors (Merkel cells and Ruffini corpuscles) detect roughness, stretch, and static stimulation on the skin. Furthermore, the fingerprint patterns and the interlocked microstructures of the human skin enhance the perception of fine texture by amplifying the vibrotactile signals during surface exploration (figure 8.1). In general, these cutaneous mechanoreceptors of the human body provide the necessary tactile information to manipulate objects with extreme accuracy (see chapter 6)

The artificial skin (eskin) was developed to mimic human skin through a combination of different materials, structures, and technologies. One of its earliest uses was in 1985,

Table 8.1
Classification of various mechanoreceptors

Classification	Pacinian corpuscle	Ruffini corpuscle	Merkel cells	Meissner's corpuscle
Adaptation rate	Fast	Slow	Slow	Fast
Effective stimuli	Temporal change in skin morphology	Vertical force detection, slippage	Spatial deformation, curvatures, edges	Temporal change in skin morphology
Sensory function	High-frequency vibration	Position, grasp, motion	Pattern detection, perception, texture	Low-frequency vibration, grip control

Source: Adapted from Dahiya 2010.




Figure 8.1

The MRs in the glabrous human skin that enable the tactile sensation. The SA MRs (*right*) respond with continuous spikes during the static stimuli, and the FA MRs (*right*) respond with spikes during the transition or the dynamic part of the stimuli.

when a flexible array with a resolution of 5 cm was attached on a robotic arm for sensing proximity (Hammock et al. 2013). Since then the nature of eskin has not changed much, as most of the eskins share similar sensors and readout characteristics along with their morphologies (Navaraj et al. 2017; Yogeswaran et al. 2018; Núñez, Manjakkal, and Dahiya 2019). Generally, they have a base substrate (bendable/stretchable) on top of which the sensing element/s (capacitive, resistive, piezoelectric, and so on) are developed. Usually, an encapsulation layer is added on top of the sensing structure to reduce the possibility of wear and tear. These devices can be bendable in order to conform to the surface of a robot's rigid body to equip them with more advanced humanlike tactile-sensing capabilities (Kappasov, Corrales, and Perdereau 2015; Yogeswaran et al. 2015; Núñez et al. 2017).

8.3 Robotic Hands with Intrinsic Tactile Sensing

Intrinsic or tightly integrated sensing, actuation, and computation elements, all embedded in 3D structures, will underpin the advances in the next generation of smart and complex systems such as humanoid robots with the capabilities to carry out cognitive tasks (Ntagios et al. 2020). The human skin is densely packed with different types of mechanoreceptors (as described in section 8.1) that support humans' ability to carry out cognitive tasks by enabling them to understand and rapidly respond to the constantly changing environment. As humans interact with the environment, the touch stimuli from these tightly coupled receptors are constantly being processed, interpreted, and stored by the brain followed by swift action from the concerned effector in response to the stimulus (Bear, Connors, and Paradiso 2020). This real-time, closed-loop interaction enables humans to respond using not only the immediate stream of information from the receptors but also the previously stored information. So for robots to be autonomous and able to carry out cognitive tasks, eskin should be able to acquire, process, and store information from the environment in a closed-loop fashion through tightly coupled sensors, actuators, and computation elements. This will enable a fast, real-time response and adaptation of the robot to its dynamic environment.

There have been some attempts toward bestowing robots with humanlike dexterity through artificial muscles, large-area eskins, computing devices, and so on, but these robots often fail to execute intricate tasks that are easily conducted by humans (Viteckova, Kutilek, and Jirina 2013; Siegart et al. 2011). The reason is that current arrangements do not explore the synergistic working of sensors, actuation, and computation to the same degree as humans. The eskins developed nowadays have some human-skin-like features, but their surface mounting comes with the challenge of wear and tear during frequent use. These issues arise from the way they are deployed on the surface of robotic bodies. The sensors need to be in direct contact with objects and often have limited protection from extreme forces and/or sharp edges. Another common problem is routing the vast amount of wires in eSkin devices to the computing unit. This often results in a potential hazard when operating a robotic system. Some of these challenges can be alleviated by embedding the sensing elements in the core structure of robots.

Additive manufacturing, or 3D printing, as it is more commonly known, has emerged over the last few decades and could offer new solutions for developing robotic parts with embedded sensors (see section 2.1). The process is based on a build sequence in which the structure is constructed from the layer-by-layer deposition of materials. As an additive method (as

opposed to a subtractive technique such as milling), 3D printing provides an ability to obtain complex 3D structures with arbitrary shapes and more geometric freedom when taking the build process into consideration. If today's single-materials-based, 3D-printing approach can be adapted to incorporate the simultaneous printing of multiple materials (e.g., plastic and metal) then there is potential for the manufacturing of "smart" objects with enhanced functionalities and with embedded sensing/electronic components (Nassar et al. 2018). This is an exciting approach for robotics because different sensing and actuating materials can be embedded into a robot's body as part of the build process. The printing of various conducting materials, along with typical plastic or polymers to create complex 3D structures, will allow the efficient use of 3D space inside these structures. The open-source nature of most fused deposition modeling (FDM) printers and their accompanying software also lends itself to widespread modifications to the printers in various ways, such as incorporating multiple printing heads, printing novel materials, and adjusting the print settings to suit a desired custom application. This being said, there are some limitations, particularly with regard to the print resolution. Nozzle diameters, build volumes, relatively slow fabrication speeds for mass production, material properties, and lack of adjustability during fabrication are some of the limiting factors of this technology. Researchers are currently working toward improving these machines via integrating other fabrication mechanisms, feedback controls, and AI (Sitthi-Amorn et al. 2015; Skylar-Scott et al. 2019). Nonetheless, the overwhelming benefits of printing rigid structural materials, soft materials, conductive inks, and sensing and actuating elements all in one fabrication method for robots in arbitrary shapes is an avenue that will spur the research in coming years.

Recent work (Ntagios et al. 2020) in which innovative hand design has been used along with multimaterial 3D printing is a good example of this approach. A 3D-printed soft capacitive sensor and associated readout electronics (e.g., a capacitance to digital converter chip on a small PCB) were embedded into the 3D-printed robotic hand (Ntagios et al. 2020). At first a five-finger 3D-printed hand was designed to have embedded actuators for movement of each of the fingers. The design consisted of multimaterial 3D printing by a modified 3D printer mounted with multiple hot ends with different materials. The hand's design was segmented into three sections: bottom, middle, and top (figure 8.2*a*). The top and bottom sections were printed with polylactic acid (PLA), a well-known 3D-printing material, and the middle part was printed with flexible thermoplastic polyurethane (TPU). In between the sections, a thin layer of acrylonitrile butadiene styrene (ABS) was printed to increase the adhesion between the sections. In this way, the entire hand was fabricated in one continuous print without the need for assembly or support material. This arrangement of materials utilized the rigidity of the PLA and ABS and the elasticity of TPU to achieve flexion of the finger joints. The hand is an underactuated and self-adapting mechanical end effector without any complex mechanical parts. This is an attractive approach to mechanical design because it achieves multiple requirements of robotic end effectors, minimizing the postprocessing and assembly time, in contrast with the more common production of robotic end effectors that utilize fabrication techniques such as machining, molding, and/or laser cutting to produce the parts of the system and are often required to implement extremely complex driving mechanisms to animate the hand (Weiner et al. 2020). Most robotic hands, especially the commercial ones, are fabricated with completely rigid materials, resulting in a massive amount of parts needing to be fabricated and assembled (Belter et al. 2013).

Further, a similar methodology was used to produce fingertips with an embedded capacitive sensor and embedded readout electronics (figure 8.2*b*). The fingertip had a simple design to enable the fabrication of the aforementioned system (figure 8.2*c*). The architecture of the phalanx imitated the structure of the human finger, with a rigid interior (bone), soft tissue, and skin. The pattern of the sensor mimicked this morphology, with a rigid PLA base and conductive and dielectric material encapsulated between the rigid PLA and the top surface made of TPU. In the core of the rigid PLA structure lay the embedded electronics. The fabrication of this part was performed in steps, the first being the printing of PLA up to the level of the two pull-up resistors, which are needed to implement the interintegrated circuit (I2C) protocol for the integrated circuit (IC) chip meant to read the capacitance variations. The subsequent steps involved the placement of resistors and the direct ink writing (DIW) of a custom-made graphite ink for interconnects. After the ink dried, a second section of PLA was printed on top until the designated area housing the PCB was mounted with a capacitance-to-digital converter IC. The PCB was placed on top, followed by further printing until complete encapsulation was reached. In the study three conductive materials, silver adhesive paste, conductive PLA, and custom-made graphite-based ink, and two dielectric materials, Ecoflex and TPU, were explored to create the capacitive sensor. Other studies have printed silicone rubber materials as part of their transducers, and they have concluded that softer materials such as Ecoflex reduce the hysteresis of the transducer (Tomo et al. 2018). Five variations of the sensor were created with a combination of these materials:




Figure 8.2

The 3D-printed hand with intrinsic tactile sensing. (a) CAD design of the hand with the smart sensing phalanx that has a soft capacitive touch sensor and an embedded readout circuit. (b) CAD design of the interior structure of the phalanx. (c) Fabrication steps for the 3D-printed phalanx.

Ecoflex-silver, TPU-silver, Ecoflex-graphite, TPU-graphite, and Ecoflex-TPU. All sensor variations were fabricated using the customized 3D printer. The conductive PLA and the TPU were deposited with fused deposition modeling (FDM) technique, and the Ecoflex was drop casted. The graphite ink was printed with direct ink writing (DIW) technique, and the silver paint was brushed, but similar techniques can also be used with the silver.

The Ecoflex-silver variation showed superior performance (figure 8.3) and a stable and repeatable response in static and dynamic conditions with a minor hysteresis effect. The superiority of the Ecoflex dielectric and the silver paste electrodes over the other devices was due to the materials' properties. The adhesion of the Ecoflex and the silver paint was found to be the strongest with respect to other samples. The silver paste, which is known to develop cracks, did not do so in the embedded configuration. This arrangement of materials and the interactions between them demonstrate an alternative approach toward sensor endurance. The embedding of sensing elements inside flexible elastomers provides the required protection to the sensing elements, thus increasing the duration of the use phase of the sensing modules and preventing costly repairs.

In recent years, a number of similar studies have been initiated that attempt to utilize this technology. Previously, most 3D-printed sensors were fabricated using the direct ink method (Muth et al. 2014). These methods were most commonly used for soft robotics and eskin-type approaches (Truby et al. 2019). Recently, more studies are using FDM techniques as well (Kaur and Kim 2019).

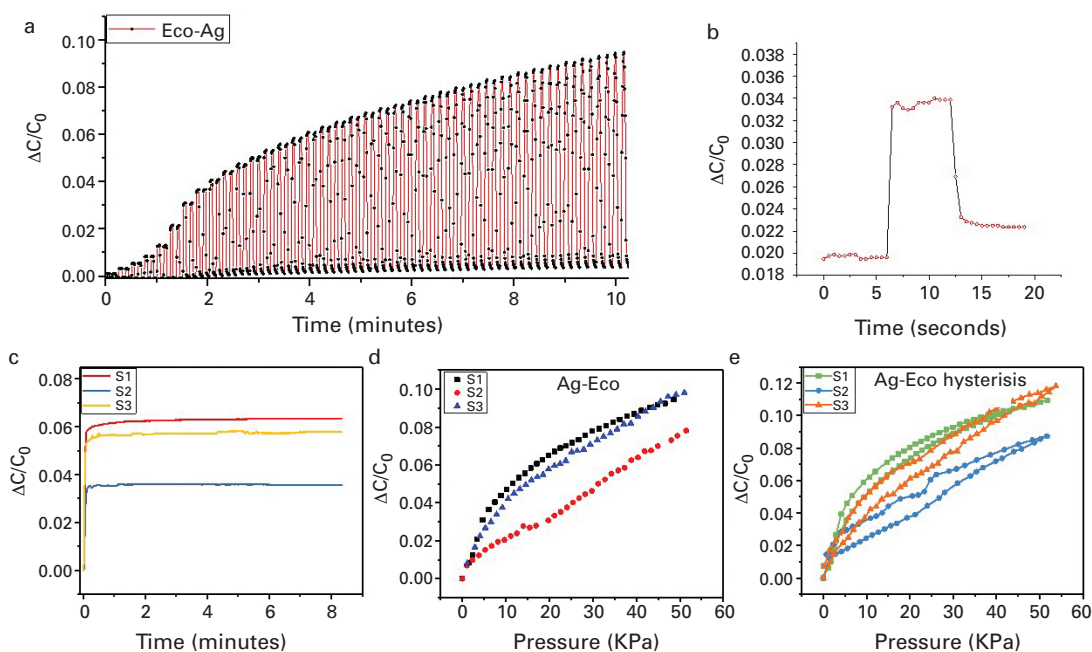


Figure 8.3

(a) Dynamic response of one of the Eco-Ag sensing devices over time with increasing pressure. (b) Relative change in capacitance of the Eco-Ag sensing device with respect to time during one of the loading-unloading cycles. (c) Response of all three sensors under constant load. (d) Relative change in capacitance with increasing pressure. (e) Hysteresis curve of the tested devices.

Obtaining complex smart structures with intrinsic sensing, actuation, and computing is the way to progress to the next era of autonomous robotic systems. The tightly integrated sensing within the 3D-printed structures could pave the way for a new generation of truly smart systems that can change their appearance and shape autonomously. In comparison with state-of-the-art robotic or prosthetic hands, this approach could lead to robust and affordable hands with more functionalities. Furthermore, the multimaterial 3D-printing methodology offers efficient use of 3D space through embedded components.

8.4 Tactile Sensor with Piezoelectric/Capacitive Stack

The dynamic and static force feedback from the skin is central to humans for daily tasks. As mentioned in 8.1, human skin contains both FA and SA mechanoreceptors. However, most of the tactile sensors reported in the literature provide either static or dynamic pressure (Yousef, Boukallel, and Althoefer 2011; Jamone et al. 2015; Kaur and Kim 2019). The spatiotemporal detection of tactile stimuli is important for texture recognition (Yousef, Boukallel, and Althoefer 2011; Dahiya et al. 2013), and for this purpose it is necessary for eskin to have the ability to detect both static and dynamic contacts. To address this issue, scientists have recently developed a new touch sensor—a stack of piezoelectric and capacitive sensors. This allows the measurement of both static and dynamic stimuli, and with the use of machine-learning or artificial intelligence (AI) tools, we can explore further cognitive skills such as detecting the texture of a curved surface (Navaraj and Dahiya 2019). This highly sensitive, capacitive-piezoelectric, flexible sensing skin with fingerprint-like patterns was formed to detect and discriminate between spatiotemporal tactile stimuli, including static and dynamic pressures and textures.

Multifunctional sensors that provide information about static and dynamic events are vital for the autonomy and dexterity of robots. In this study, to compensate for the inability of the piezoelectric sensor to perform static sensing, an integrated capacitive sensor was introduced. Thus, a capacitive-piezoelectric sensor stack was formed to mimic human skin's SA and FA mechanoreceptors (figure 8.4). The sensor was encapsulated within the 3D-printed distal phalanx of the index finger, using fingertip patterns from TPU. This pattern enhanced the detection capability of the system to identify surface roughness. This is a significant leap forward, as most of the surface roughness systems developed prior to this work have relied heavily on large area arrays (Drimus et al. 2014; Lee, Kukreja, and Thakor 2017).

The tactile sensor had a floating electrode-based capacitive structure in tandem with a piezoelectric structure. The sensor utilized two soft elastomers with low and high Young's modulus. This arrangement enabled high sensitivity at low pressures, due to the softer elastomer, without saturating at higher pressures, due to the high Young's modulus elastomer. At static pressure, the elastomers compressed, and the floating electrode moved closer to the signal and ground electrodes (figure 8.4). The sensor stack was integrated into the distal phalanx of the index finger of a 3D-printed prosthetic/robotic hand. The sensing device was covered with fingerprint ridges made from TPU polymer 3D-printed filament. The ridges were positioned in a staggered fashion to provide robust protection, in a way similar to human skin.

Early studies in this field have implemented classifiers with tactile sensors utilizing Fourier transform wavelets. Researchers have concluded that a change in texture over time

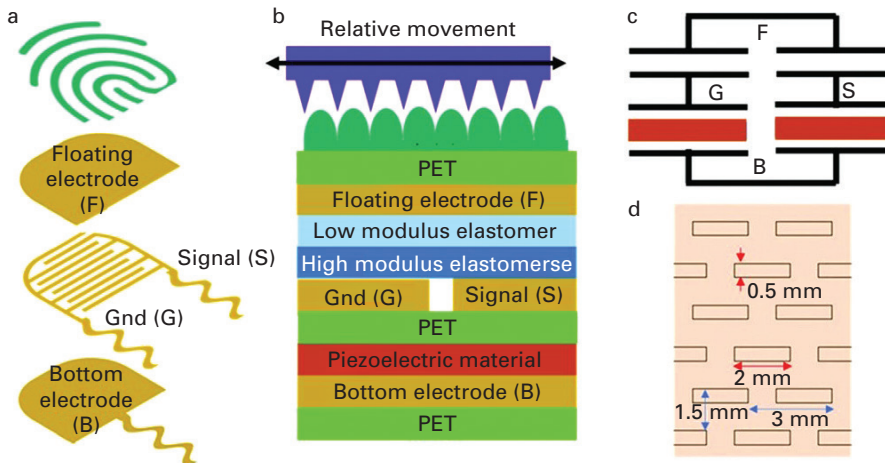


Figure 8.4

(a) Schematic illustration of the biomimetic sensory stack. (b) The layers in the sensory stack, with fingerprint ridges shown at the top. (c) An equivalent diagram of the biomimetic sensory stack. (d) The dimensions of the designed fingerprint ridges realized via 3D printing using NinjaFlex.

is an important factor between surfaces with irregular textures. A short-term Fourier transform could be used to explore more irregular surfaces (Jamali and Sammut 2011).

In Navaraj and Dahiya (2019), a biologically plausible wavelet transform was used to encode the sensor's output into spike trains based on a leaky integrated-and-fire (LIF) model. The spikes were classified with a tempotron classifier using a biological observed spike timing-dependent plasticity (STDP) mechanism learning algorithm. With this approach nonplanar texture surfaces can be classified, unlike prior works. This was made possible with a six-degrees-of-freedom robotic arm that maintained constant static pressure on the surface of the object. The data were fed to a wavelet-based processing algorithm, using the Gabor wavelet transform (GWT) instead of the common Fourier transform. This approach offers localization in time and frequency domain, and at the same time wavelet transform appears to be a more plausible approach in biological systems. To further prove the point, the results were also presented using short-time Fourier transform (STFT) with a window size of one hundred samples. After the GWT transform, the data were encoded into latency-coded spike trains, as this is the assumed reason why biological systems have such a fast response to dynamic stimuli. An LIF model was used for the spike model, while the amplitude represented how fast the spike was elicited within the time span. This work was tested to prove whether textures can be perceived with a single biomimetic sensory stack. To prove the truthfulness of the above statement, hook-and-loop fasteners were used as textures for binary classification. The classification was conducted using both planar and nonplanar surfaces to remove possible biases. One hundred planar scans, fifty concave and fifty convex, were recorded, with each scan comprising both hook and loop textures. Training data consisted of 160 randomly selected samples for the neural network and 40 for testing.

Figures 8.5*a–d* show the system's response and easily demonstrate that the loops produced higher amplitude signals than the hooks due to loops interacting more with the fingertip patterns. In the training error over the number of epochs, it is also clear that the

GWT approach to texture recognition is superior to the traditional STFT method. The STFT-based approach has an accuracy of 95.3 percent, while the GWT-based approach offers 99.45 percent accuracy for the same window of time (figure 8.5e). In conclusion, the output of the sensory stack under a closed-loop system was able to classify textures with a maximum accuracy of 99.45 percent, which also demonstrated the possibility that a single sensory stack may be sufficient for texture classification.

8.5 Integrated Sensing and Actuation Technology

This section examines the research focused on integrating sensors and actuators for an advanced eSkin. To utilize the full potential of robots, it is important to enable them to interact with dexterity and cognitive capabilities, as well as learn from their resulting interaction with the environment. The purposeful employment of a robot's environment is proposed in the context of developmental robotics in section 6.3. Future robots should be able to deal with the uncertainty of the natural environment by continually learning, reasoning, and sharing their knowledge. As previously discussed in section 8.2, eskin is one of the effective approaches that researchers have used to achieve this. However, existing robots are mostly equipped with eskin having only sensing capabilities. As mentioned in

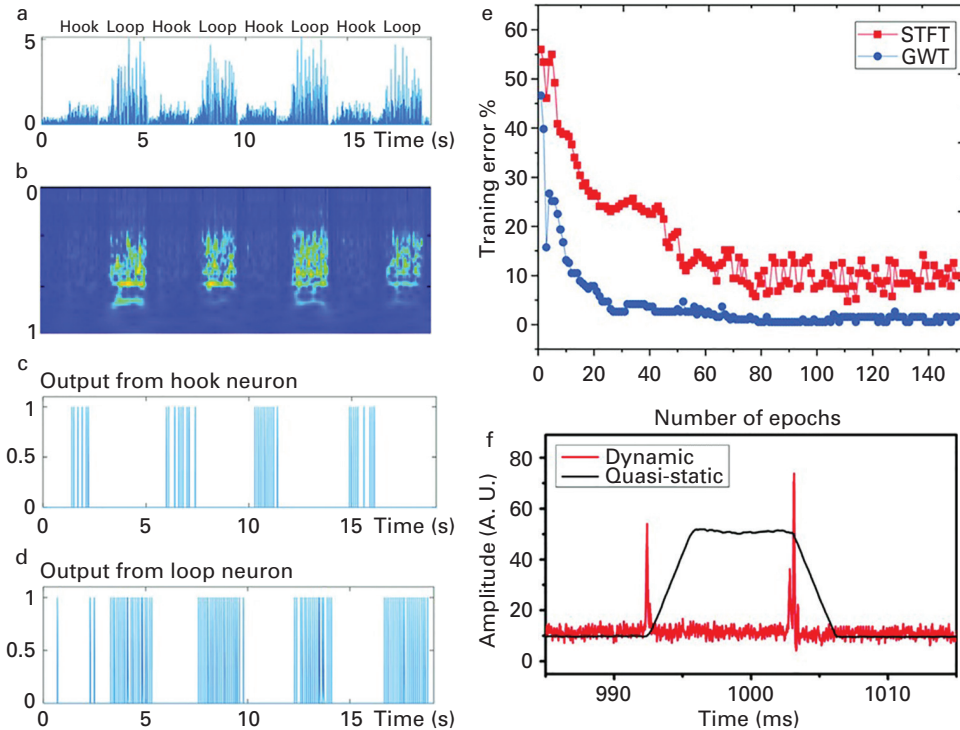


Figure 8.5

(a) A typical recorded signal from the dynamic scan. (b) Gabor wavelet scalogram. (c, d) Output from the tempotron classifier neuron corresponding to the (c) hook and (d) loop. (e) Training error versus number of epochs comparing STFT and GWT-based features. (f) Plot of the wirelessly transmitted live data acquired via the rqt plot of the ROS package.

section 8.2, researchers have designed tactile sensors for eskin using various material (Yogeswaran et al. 2015) structures (Mannsfeld et al. 2010; Gong et al. 2014; Wang et al. 2014), morphologies (Dahiya and Gori 2010; Navaraj et al. 2017; Navaraj and Dahiya 2019), and transduction methods (Dahiya et al. 2011; Adami et al. 2012; Dahiya and Valle 2013; Khan et al. 2015; Gupta, Shakthivel, et al. 2018; Gupta, Yogeswaran, et al. 2018; Hannah et al. 2018; Kawasetsu et al. 2018; Yogeswaran et al. 2018), with some mimicking the human skin's features, such as fingertip-like patterns on the surface and integrated static and dynamic sensors (Navaraj and Dahiya 2019). However, the complexity of eskin goes beyond just integrating various types of touch sensors on flexible substrates (Núñez et al. 2017).

Seamless integration of both sensing and actuation capabilities will improve the granularity of haptic information inherent in the next generation of eskin (Dahiya et al. 2019), enabling a substantial contribution to AI systems. Robots donned in such eskin will have humanlike dexterity, cognitive skills, and physical abilities, as they will be able to learn from their environment via rich and diverse information. In this context, some studies have explored adding sensing capabilities to different types of actuators to obtain information regarding the degree of displacement produced during actuation. These actuators include electromagnetic (Andò and Marletta 2016; Do et al. 2018), pneumatic (Yeo et al. 2016), and electroactive polymers (EAPs; Nakamura and Kawakami 2019) with different operating principles and materials (Chen et al. 2019). Unlike electromagnetic actuators and ionic EAPs (Asaka et al. 2013), the majority are unable to provide bidirectional actuation, vibrotactile feedback, or a high level of displacement due to limitations in the actuation principle and/or materials used (Biswas and Visell 2019). Further, the majority of the actuators with integrated sensing functions are manufactured either on paper (Phan et al. 2017; Amjadi and Sitti 2018) or with EAPs (Jung, Kim, and Choi 2008) that require relatively high voltages (~ 150 V per micrometer displacement; Yeo et al. 2016). Electromagnetic actuators are capable of providing high displacement (up to 1 mm and a high force ~ 5 mN/mm² at 5 V; Guo et al. 2018; Noguchi, Nagai, and Kawamura 2018) as well as bidirectional actuation (Bintoro et al. 2005) and vibrotactile feedback (Do et al. 2018) at different frequencies (≤ 1 Hz and > 500 Hz). In particular, bidirectional actuation is advantageous in the manipulation of the direction of actuation, as it provides options for controlled multidirectional displacement (Cho and Ahn 2002). In the effort to make electromagnetic actuators intrinsically soft and wearable, the field of magnetoelectronics has also been rapidly gaining attention (Hellebrekers, Kroemer, and Majidi 2019). In this case, flexible magnets and elastomers mixed with ferromagnetic materials are harnessed for the purpose of actuation (Almansouri et al. 2019; Hintze et al. 2014). However, electromagnetic actuators have so far been those most employed for actuation purposes (Said et al. 2016; Paknahad and Tahmasebipour 2019), with primary applications in micropumps (Said et al. 2018) and tactile displays (Zárate and Shea 2016). By integrating sensing capabilities in electromagnetic actuators, multidirectional actuation capability and excellent controllability could be adequately harnessed to advance applications in the realization of soft eskin with both sensing and haptic feedback capabilities.

Electromagnetic actuators (EMAs) function by converting magnetic energy into mechanical energy and are generally governed by three fundamental laws: the Lorentz law, Faraday's law, and the Biot-Savart law (Gomis-Bellmunt and Campanile 2009). Actuation in

EMAs occurs through the interaction of the magnetic field (produced by a current through a coil) with a permanent magnet and/or a ferromagnetic material (Kawasetu et al. 2018). This interaction produces either a repulsive or an attractive force applied directly to a membrane or plunger, thereby causing displacement. This repulsive or attractive force is utilized to achieve the repulsion and attraction of soft membranes of the eskin. Principally, electromagnetic actuation occurs by means of two main circuits: 1) the electrical circuit that establishes the current and voltages and 2) the magnetic circuit that establishes the magnetic field strength and flux. The current I produces the controllable magnetic field \vec{H} , while the magnetic field produces the magnetic flux ϕ and the magnetic flux density \vec{B} (equation 8.1).

$$\vec{B} = \mu_r \mu_0 \vec{H}, \quad (8.1)$$

where μ_r = relative permeability of the material, μ_0 = permeability of the vacuum, and the magnetic constant = $4\pi \times 10^{-7} H$.

The addition of intrinsic sensing to electromagnetic actuators is advantageous, as mentioned in section 8.3 and shown in research (Ozioko, Navaraj, et al. 2018; Ozioko, Hersh, and Dahiya 2018, 2019; Ozioko, Karipoth, et al. 2021). Figure 8.6a shows this principle and the device structure composed of a tactile sensing (piezoresistive) layer integrated on top of a permanent magnet that is part of a flexible electromagnetic, coil-based actuator. The device can detect contact force via the piezoresistive layer and simultaneously produce a proportional actuation using the electromagnetic actuator. Figure 8.6a shows the device before actuation, while figure 8.6b and figure 8.6c show the device during different actuation modes. During actuation, the top layer is attracted to or repelled by the coil, as shown in figure 8.6b and figure 8.6c respectively, in accordance with the electromagnetic principle previously described in this section. In each case, there are two possible states, the vibration state and the nonvibration state, depending on the direction of the supplied current. When a constant current is supplied, the device operates in a nonvibration state. The vibration state occurs when the pulsating current of a given frequency is supplied through the coil. This makes it possible to control the speed, movement, and direction of the top layer via the manipulation of the magnitude and direction of the supplied current. Hence, eskin with this feature can be controllably tuned as required.

Figure 8.7 shows a more detailed operating principle of the device. The two main modules of the device (sensing and actuation) are controlled by the sensing and actuation module, respectively. Figure 8.7a shows the soft, piezoresistive sensing layer. This sensing layer could be realized using any soft sensing layer, but in this case a graphite ink was encapsulated using Sil-Poxy™. When an external force is applied to the sensing layer, the particles of graphite move closer to one another from distance $d1$ and $d2$ to $d1 + \Delta d1$ and $d2 + \Delta d2$, respectively. This creates a closer conducting network that causes a reduction in resistance of the material from R (figure 8.7a1) to $R + \Delta R$ (figure 8.7a2). Figure 8.7c shows what happens when external pressure is applied to the sensing layer. In this case, a change in resistance (ΔR) occurs as read by the sensing control module. This resistance shift causes a change in current (ΔI) flowing through the spiral coil that is driven by the actuation control module. This change in current in turn causes a proportional change in the magnetic field produced by the spiral coil that leads to a change in the force of actuation. This change in actuation force causes the top layer to move away from the coil due to repulsion or closer to the coil due to attraction. Therefore, this device takes advantage

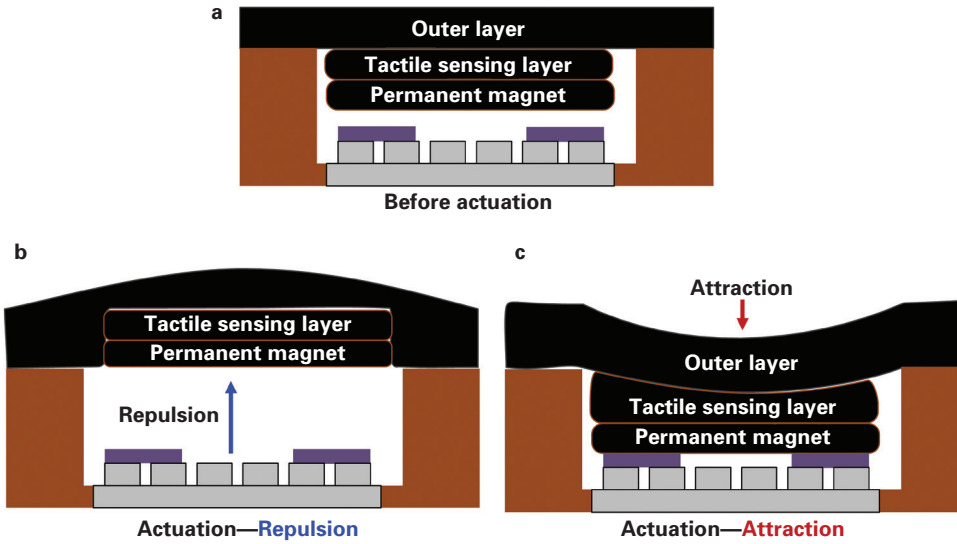


Figure 8.6 Actuation modes. (a) Before actuation. (b) Repulsion mechanism. (c) Attraction mechanism.

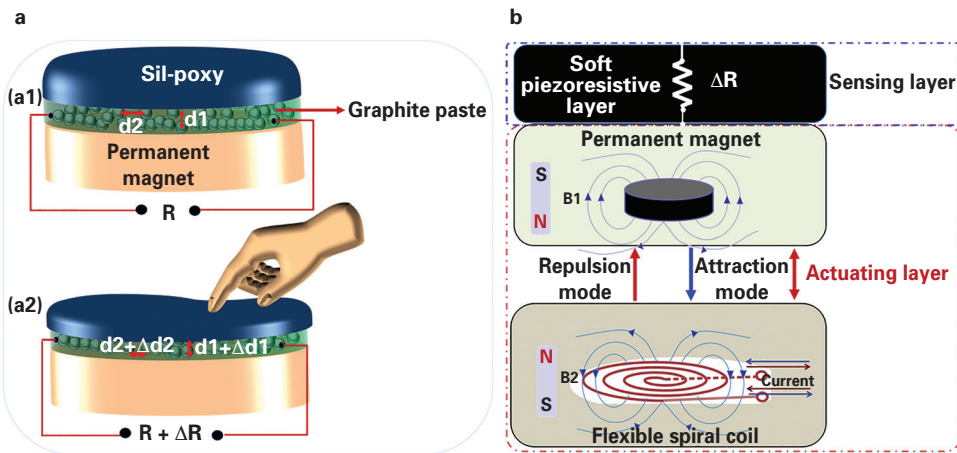


Figure 8.7 (a) The structure and principle of the piezoresistive layer. (b) The structure and principle of the actuation layer.

of the sensing ability of the piezoresistive layer and the magnetic interaction between the coil and the permanent magnet to produce simultaneous sensing and actuation. Additionally, the sensing and actuation could be independently controlled using digital logic gates and a microcontroller programmed with corresponding algorithms.

Figure 8.8 shows the response of the sensing layer of the integrated device alone as well as that of simultaneous sensing and actuation. This result illustrates that the self-controllability characteristic of the integrated device makes the concept advantageous for use in future tunable eskin, enabling controllability and the extraction of richer information. Future applications could explore embedding this integrated device in a robotic fingertip to control its stiffness for improved grasping of objects.

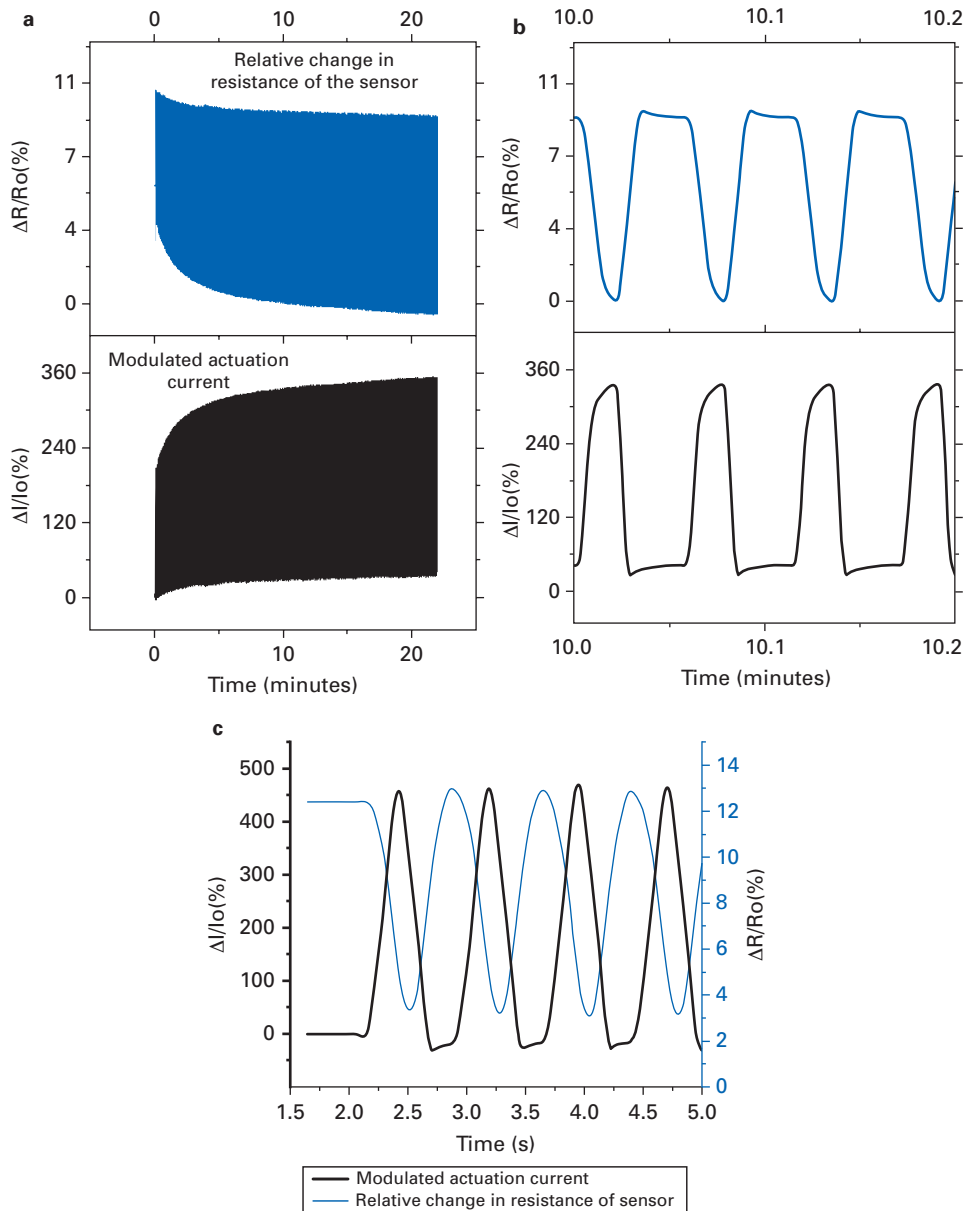


Figure 8.8

(a) Output of sensing layer and the modulated current through the actuating coil during simultaneous sensing and actuation at approximately 0.25 Hz for approximately twenty-two minutes. (b) Zoom-in of figure 8.8a demonstrating that as the resistance decreases, the current through the actuator increases. It also shows the stability of the device for approximately twenty-two minutes of continuous use. (c) Change in current and resistance during simultaneous sensing and actuation at 1 Hz.

8.6 Conclusion

This chapter presented the current state and development of tactile sensing and actuation technologies in robotic skin along with new approaches toward biomimetic and bioinspired tactile sensing and computing. Current fabrication techniques and their limitations and drawbacks were discussed. The growth of 3D printing and the advantages it provides were examined, as well as how this new technology can enhance current tactile and actuation systems. The two tactile sensor structures presented in sections 8.3 and 8.4 do in many ways mimic human skin's functionality. The 3D-printed hand with intrinsic tactile sensing discussed in section 8.2 has the embedded actuation, capacitive sensors in the distal phalanges, and embedded electronics capable of reading the capacitive value and transmitting the digital information to the microcontroller. In this way the wear and tear issue of eskin is alleviated, along with the wiring complexity issue. The biomimetic tactile sensor presented in section 8.3 uses a sensory stack to simultaneously measure dynamic and static conditions. Data from sensors were fed into a neural network that could classify two textured surfaces with an extreme accuracy rate of 99.45 percent. In the study, a comparison between the commonly used short-time Fourier transform and a biomimetic Gabor wavelet transform was performed to explore the superior system. The piezoresistive sensor with integrated actuation presented in section 8.5 can provide vibrotactile feedback. Devices such as these have the potential to advance soft robotics by allowing such robots to “squeeze” while continuing to sense ambient conditions.

In general, the case studies presented show how eskin research is advancing toward tightly coupled sensing, actuation, and computation. The key cognitive skills needed to advance robot capabilities include memory, decision-making, action understanding, and prediction. The technologies discussed in this chapter open opportunities for achieving these skills by allowing robots to effectively sense their environment and process, store, and use the obtained information to respond to their dynamic environment. This can have a significant impact on human-robot interaction—for instance, humans are able to extract important information from tactile stimuli that depends not only on the underlying touch characteristics but also on the context of the touch, culture, and emotions of the individuals who are communicating. So enabling robots not only to sense tactile information but also to understand the intended meaning of touch has great potential to advance robot cognition as well as human-robot interaction.

Additional Reading and Resources

- This edited volume provides a complete overview of tactile sensing in humans. It includes definitions and classification. It also classifies all transduction methods to realize tactile sensors and materials. Dahiya, Ravinder S., and Maurizio Valle. 2013. *Robotic Tactile Sensing: Technologies and System*. Berlin: Springer Science and Business Media.
- A compact volume conveying a great deal of information on sensing and actuation. This volume provides information on sensing for broad variety of stimuli. Extensive description is given to robot motion, both for soft and rigid robots, tackling some control algorithms. Siciliano, Bruno, and Oussama Khatib, eds. 2016. *Springer Handbook of Robotics*. Berlin: Springer.

- A special issue presenting the latest work on flexible electronics and eskin. Dahiya, Ravinder, Deji Akinwande, and Joseph S. Chang. 2019. “Flexible Electronic Skin: From Humanoids to Humans.” Special Issue, *Proceedings of the IEEE* 107 (10): 2011–2015.
- Basic knowledge of different types of tactile-sensing mechanisms for nonexperts. Explains the different stimuli and basic circuitry used for reading the outputs: <https://www.elprocus.com/tactile-sensor-types-and-its-working/>.
- BEST (Bendable Electronics and Sensing Technologies) YouTube channel, containing robotic/prosthetic videos with tactile sensing, 3D printing, and more: <https://www.youtube.com/channel/UCOODG132wFmWSTPPBUARAvA/>.

References

- Adami, Andrea, Ravinder S. Dahiya, Cristian Collini, Davide Cattin, and Leandro Lorenzelli. 2012. “POSFET Touch Sensor with CMOS Integrated Signal Conditioning Electronics.” *Sensors and Actuators A: Physical* 188:75–81.
- Almansouri, Abdullah S., Nouf A. Alsharif, Mohammed A. Khan, Liam Swanepoel, Altnayn Kaidarova, Khaled N. Salama, and Jurgen Kosel. 2019. “An Imperceptible Magnetic Skin.” *Advanced Materials Technologies* 4 (10): 1900493.
- Amjadi, Morteza, and Metin Sitti. 2018. “Self-Sensing Paper Actuators Based on Graphite-Carbon Nanotube Hybrid Films.” *Advanced Science* 5 (7): 1800239.
- Andò, Bruno, and Vincenzo Marletta. 2016. “An All-Inkjet Printed Bending Actuator with Embedded Sensing Feature and an Electromagnetic Driving Mechanism.” *Actuators* 5 (3): 21.
- Argall, Brenna D., and Aude G. Billard. 2010. “A Survey of Tactile Human-Robot Interactions.” *Robotics and Autonomous Systems* 58 (10): 1159–1176.
- Asaka, Kinji, Ken Mukai, Takushi Sugino, and Kenji Kiyohara. 2013. “Ionic Electroactive Polymer Actuators Based on Nano-carbon Electrodes.” *Polymer International* 62 (9): 1263–1270.
- Bear, Mark, Barry Connors, and Michael Paradiso. 2020. *Neuroscience: Exploring the Brain*. Burlington, MA: Jones and Bartlett.
- Belter, Joseph T., Jacob L. Segil, Aaron M. Dollar, and Richard F. Weir. 2013. “Mechanical Design and Performance Specifications of Anthropomorphic Prosthetic Hands: A Review.” *Journal of Rehabilitation Research and Development* 50 (5): 599.
- Bintoro, J. S., A. D. Papania, Y. H. Berthelot, and P. J. Hesketh. 2005. “Bidirectional Electromagnetic Microactuator with Microcoil Fabricated on a Single Wafer: Static Characteristics of Membrane Displacements.” *Journal of Micromechanics and Microengineering* 15 (8): 1378.
- Biswas, Shantonu, and Yon Visell. 2019. “Emerging Material Technologies for Haptics.” *Advanced Materials Technologies* 4 (4): 1900042.
- Chen, Luzhuo, Mingcen Weng, Peidi Zhou, Feng Huang, Changhong Liu, Shoushan Fan, and Wei Zhang. 2019. “Graphene-Based Actuator with Integrated-Sensing Function.” *Advanced Functional Materials* 29 (5): 1806057.
- Cho, Hyoung J., and Chong H. Ahn. 2002. “A Bidirectional Magnetic Microactuator Using Electroplated Permanent Magnet Arrays.” *Journal of Microelectromechanical Systems* 11 (1): 78–84.
- Dahiya, Ravinder S. 2019. “E-Skin: From Humanoids to Humans.” *Proceedings of the IEEE* 107 (2): 247–252.
- Dahiya, Ravinder S., Deji Akinwande, and Joseph S. Chang. 2019. “Flexible Electronic Skin: From Humanoids to Humans [Scanning the Issue].” *Proceedings of the IEEE* 107 (10): 2011–2015.
- Dahiya, Ravinder S., Davide Cattin, Andrea Adami, Cristian Collini, Leonardo Barboni, Maurizio Valle, Leandro Lorenzelli, Roberto Oboe, Giorgio Metta, and Francesca Brunetti. 2011. “Towards Tactile Sensing System on Chip for Robotic Applications.” *IEEE Sensors Journal* 11 (12): 3216–3226.
- Dahiya, Ravinder S., and Monica Gori. 2010. “Probing with and into Fingerprints.” *Journal of Neurophysiology* 104 (1): 1–3.
- Dahiya, Ravinder S., Giorgio Metta, Maurizio Valle, and Giulio Sandini. 2010. “Tactile Sensing—from Humans to Humanoids.” *IEEE Transactions on Robotics* 26 (1): 1–20.
- Dahiya, Ravinder S., Philipp Mittendorfer, Maurizio Valle, Gordon Cheng, and Vladimir J. Lumelsky. 2013. “Directions toward Effective Utilization of Tactile Skin: A Review.” *IEEE Sensors Journal* 13 (11): 4121–4138.

- Dahiya, Ravinder S., and Maurizio Valle. 2013. *Robotic Tactile Sensing: Technologies and System*. Berlin: Springer Science and Business Media.
- Dahiya, Ravinder S., Nivasan Yogeswaran, Fengyuan Liu, Libu Manjakkal, Etienne Burdet, Vincent Hayward, and Henrik Jörintell. 2019. "Large-Area Soft E-Skin: The Challenges Beyond Sensor Designs." *Proceedings of the IEEE* 107 (10): 2016–2033.
- Decherchi, Sergio, Paolo Gastaldo, Ravinder S. Dahiya, Maurizio Valle, and Rodolfo Zunino. 2011. "Tactile-Data Classification of Contact Materials Using Computational Intelligence." *IEEE Transactions on Robotics* 27 (3): 635–639.
- Do, Thanh Nho, Hung Phan, Thuc-Quyen Nguyen, and Yon Visell. 2018. "Miniature Soft Electromagnetic Actuators for Robotic Applications." *Advanced Functional Materials* 28 (18): 1800244.
- Drimus, Alin, Gert Kootstra, Arne Bilberg, and Danica Kragic. 2014. "Design of a Flexible Tactile Sensor for Classification of Rigid and Deformable Objects." *Robotics and Autonomous Systems* 62 (1): 3–15.
- Gomis-Bellmunt, Oriol, and Lucio Flavio Campanile. 2009. "Design Rules for Actuators in Active Mechanical Systems." Berlin: Springer Science and Business Media.
- Gong, Shu, Willem Schwalb, Yongwei Wang, Yi Chen, Yue Tang, Jye Si, Bijan Shirinzadeh, and Wenlong Cheng. 2014. "A Wearable and Highly Sensitive Pressure Sensor with Ultrathin Gold Nanowires." *Nature Communications* 5 (1): 1–8.
- Guo, Rui, Lei Sheng, HengYi Gong, and Jing Liu. 2018. "Liquid Metal Spiral Coil Enabled Soft Electromagnetic Actuator." *Science China Technological Sciences* 61 (4): 516–521.
- Gupta, Shoubhik, William T. Navaraj, Leandro Lorenzelli, and Ravinder S. Dahiya. 2018. "Ultra-thin Chips for High-Performance Flexible Electronics." *npj Flexible Electronics* 2 (1): 1–17.
- Gupta, Shoubhik, Dhayalan Shakthivel, Leandro Lorenzelli, and Ravinder S. Dahiya. 2018. "Temperature Compensated Tactile Sensing Using MOSFET with P(VDF-Trfe)/Batio3 Capacitor as Extended Gate." *IEEE Sensors Journal* 19 (2): 435–442.
- Gupta, Shoubhik, Nivasan Yogeswaran, Flavio Giacomozzi, Leandro Lorenzelli, and Ravinder S. Dahiya. 2018. "Flexible AlN Coupled MOSFET Device for Touch Sensing." In *2018 IEEE Sensors*, 1–4. New York: IEEE. <https://doi.org/10.1109/ICSENS.2018.8589628>.
- Hammock, Mallory L., Alex Chortos, Benjamin C-K. Tee, Jeffrey B-H. Tok, and Zhenan Bao. 2013. "25th Anniversary Article: The Evolution of Electronic Skin (E-Skin): A Brief History, Design Considerations, and Recent Progress." *Advanced Materials* 25 (42): 5997–6038.
- Hannah, Stuart, Alan Davidson, Ivan Glesk, Deepak Uttamchandani, Ravinder S. Dahiya, and Helena Gleskova. 2018. "Multifunctional Sensor Based on Organic Field-Effect Transistor and Ferroelectric Poly (Vinylidene Fluoride Trifluoroethylene)." *Organic Electronics* 56:170–177.
- Hellebrekers, Tess, Oliver Kroemer, and Carmel Majidi. 2019. "Soft Magnetic Skin for Continuous Deformation Sensing." *Advanced Intelligent Systems* 1 (4): 1900025.
- Hintze, C., D. Yu Borin, D. Ivaneyko, V. Toshchevikov, M. Saphiannikova-Grenzer, and G. Heinrich. 2014. "Soft Magnetic Elastomers with Controllable Stiffness: Experiments and Modelling." *Kgk-Kautschuk Gummi Kunststoffe* 67 (4): 53–59.
- Jamali, Nawid, and Claude Sammut. 2011. "Majority Voting: Material Classification by Tactile Sensing Using Surface Texture." *IEEE Transactions on Robotics* 27 (3): 508–521.
- Jamone, Lorenzo, Lorenzo Natale, Giorgio Metta, and Giulio Sandini. 2015. "Highly Sensitive Soft Tactile Sensors for an Anthropomorphic Robotic Hand." *IEEE Sensors Journal* 15 (8): 4226–4233.
- Jung, Kwangmok, Kwang J. Kim, and Hyouk Ryeol Choi. 2008. "A Self-Sensing Dielectric Elastomer Actuator." *Sensors and Actuators A: Physical* 143 (2): 343–351.
- Kappasov, Zhanat, Juan-Antonio Corrales, and Véronique Perdereau. 2015. "Tactile Sensing in Dexterous Robot Hands." *Robotics and Autonomous Systems* 74:195–220.
- Kaur, Manpreet, and Woo Soo Kim. 2019. "Toward a Smart Compliant Robotic Gripper Equipped with 3D-Designed Cellular Fingers." *Advanced Intelligent Systems* 1 (3): 1900019.
- Kawasetsu, Takumi, Takato Horii, Hisashi Ishihara, and Minoru Asada. 2018. "Flexible Tri-axis Tactile Sensor Using Spiral Inductor and Magnetorheological Elastomer." *IEEE Sensors Journal* 18 (14): 5834–5841.
- Khan, Saleem, Wenting Dang, Leandro Lorenzelli, and Ravinder S. Dahiya. 2015. "Flexible Pressure Sensors Based on Screen-Printed P(VDF-TrFE) and P(VDF-TrFE)/MWCNTs." *IEEE Transactions on Semiconductor Manufacturing* 28 (4): 486–493.
- Kumaresan, Y., O. Ozioko, and Ravinder S. Dahiya. 2021. "Multifunctional Electronic Skin with a stack of Temperature and Pressure Sensor Arrays." *IEEE Sensors Journal*. doi:10.1109/jsen.2021.3055458.
- Lee, Wang Wei, Sunil L. Kukreja, and Nitish V. Thakor. 2017. "Discrimination of Dynamic Tactile Contact by Temporally Precise Event Sensing in Spiking Neuromorphic Networks." *Frontiers in Neuroscience* 11:5.

- Luo, Shan, Joao Bimbo, Ravinder S. Dahiya, and Hongbin Liu. 2017. "Robotic Tactile Perception of Object Properties: A Review." *Mechatronics* 48:54–67.
- Mannsfeld, Stefan C. B., Benjamin C. K. Tee, Randall M. Stoltenberg, Christopher V. H. H. Chen, Soumendra Barman, Beinn V. O. Muir, Anatoliy N. Sokolov, Colin Reese, and Zhenan Bao. 2010. "Highly Sensitive Flexible Pressure Sensors with Microstructured Rubber Dielectric Layers." *Nature Materials* 9 (10): 859–864.
- Mohd Said, Muzalifah, Jumril Yunas, Badariah Bais, Azrul Azlan Hamzah, and Burhanuddin Yeop Majlis. 2018. "The Design, Fabrication, and Testing of an Electromagnetic Micropump with a Matrix-Patterned Magnetic Polymer Composite Actuator Membrane." *Micromachines* 9 (1): 13.
- Muth, Joseph T., Daniel M. Vogt, Ryan L. Truby, Yiğit Mengüç, David B. Kolesky, Robert J. Wood, and Jennifer A. Lewis. 2014. "Embedded 3D Printing of Strain Sensors within Highly Stretchable Elastomers." *Advanced Materials* 26 (36): 6307–6312.
- Nakamura, Atsushi, and Shotaro Kawakami. 2019. "An Actuator-Sensor Hybrid Device Made of Carbon-Based Polymer Composite for Self-Sensing Systems." *AIP Advances* 9 (6): 065311.
- Nassar, Habib, Markellos Ntagios, William T. Navaraj, and Ravinder S. Dahiya. 2018. "Multi-Material 3D Printed Bendable Smart Sensing Structures." In *2018 IEEE Sensors*, 1–4. New York: IEEE. <https://doi.org/10.1109/ICSENS.2018.8589625>.
- Navaraj, William T., and Ravinder S. Dahiya. 2019. "Fingerprint-Enhanced Capacitive-Piezoelectric Flexible Sensing Skin to Discriminate Static and Dynamic Tactile Stimuli." *Advanced Intelligent Systems* 1 (7): 1900051.
- Navaraj, William T., Carlos García Núñez, Dhayalan Shakthivel, Vincenzo Vinciguerra, Fabrice Labeau, Duncan H. Gregory, and Ravinder S. Dahiya. 2017. "Nanowire FET Based Neural Element for Robotic Tactile Sensing Skin." *Frontiers in Neuroscience* 11:501.
- Noguchi, Takuya, Sakahisa Nagai, and Atsuo Kawamura. 2018. "Electromagnetic Linear Actuator Providing High Force Density per Unit Area without Position Sensor as a Tactile Cell." *IEEE Journal of Industry Applications* 7 (3): 259–265.
- Ntagios, Markellos, Habib Nassar, Abhilash Pullanchiyodan, William T. Navaraj, and Ravinder S. Dahiya. 2020. "Robotic Hands with Intrinsic Tactile Sensing via 3D Printed Soft Pressure Sensors." *Advanced Intelligent Systems* 2 (6): 1900080.
- Núñez, Carlos García, Libu Manjakkal, and Ravinder S. Dahiya. 2019. "Energy Autonomous Electronic Skin." *npj Flexible Electronics* 3 (1): 1–24.
- Núñez, Carlos García, William T. Navaraj, Emre O. Polat, and Ravinder S. Dahiya. 2017. "Energy-Autonomous, Flexible, and Transparent Tactile Skin." *Advanced Functional Materials* 27 (18): 1606287.
- Ozioko, Oliver, Marion Hersh, and Ravinder S. Dahiya. 2018. "Inductance-Based Flexible Pressure Sensor for Assistive Gloves." In *2018 IEEE Sensors*, 1–4. New York: IEEE. <https://doi.org/10.1109/ICSENS.2018.8589826>.
- Ozioko, Oliver, Marion Hersh, and Ravinder S. Dahiya. 2019. "Inductance-Based Soft and Flexible Pressure Sensors Using Various Compositions of Iron Particles." In *2019 IEEE Sensors*, 1–4. New York: IEEE. <https://doi.org/10.1109/SENSORS43011.2019.8956646>.
- Ozioko, Oliver, Prakash Karipath, P. Escobedo, Markellos Ntagios, A. Pullanchiyodan, and Ravinder S. Dahiya. 2021. "SensAct: The Soft and Squishy Tactile Sensor with Integrated Flexible Actuator." *Advanced Intelligent Systems* 3 (3): 1900145.
- Ozioko, Oliver, Prakash Karipath, Marion Hersh, and Ravinder S. Dahiya. 2020. "Wearable Assistive Tactile Communication Interface Based on Integrated Touch Sensors and Actuators." *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 28 (6): 1344–1352.
- Ozioko, Oliver, William T. Navaraj, Marion Hersh, and Ravinder S. Dahiya. 2020. "Tacsac: A Wearable Haptic Device with Capacitive Touch-Sensing Capability for Tactile Display." *Sensors* 20:4780.
- Ozioko, Oliver, William T. Navaraj, Nivasan Yogeswaran, Marion Hersh, and Ravinder S. Dahiya. 2018. "Tactile Communication System for the Interaction between Deafblind and Robots." In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication*, 416–421. New York: IEEE.
- Paknahad, Ali Asghar, and Mohammad Tahmasebipour. 2019. "An Electromagnetic Micro-actuator with PDMS-Fe₃O₄ Nanocomposite Magnetic Membrane." *Microelectronic Engineering* 216:111031.
- Phan, Hoang-Phuong, Toan Dinh, Tuan-Khoa Nguyen, Ashkan Vatani, Abu Riduan Md Faisal, Afzaal Qamar, Atieh Ranjbar Kermany, Dzung Viet Dao, and Nam-Trung Nguyen. 2017. "Self-Sensing Paper-Based Actuators Employing Ferromagnetic Nanoparticles and Graphite." *Applied Physics Letters* 110 (14): 144101.
- Roche, Ellen T., Robert Wohlfarth, Johannes T. B. Overvelde, Nikolay V. Vasilyev, Frank A. Pigula, David J. Mooney, Katia Bertoldi, and Conor J. Walsh. 2014. "A Bioinspired Soft Actuated Material." *Advanced Materials* 26 (8): 1200–1206.
- Said, Muzalifah Mohd, Jumril Yunas, Roer Eka Pawinanto, Burhanuddin Yeop Majlis, and Badariah Bais. 2016. "PDMS Based Electromagnetic Actuator Membrane with Embedded Magnetic Particles in Polymer Composite." *Sensors and Actuators A: Physical* 245:85–96.

- Siegwart, Roland, Illah Reza Nourbakhsh, and Davide Scaramuzza. 2011. *Introduction to Autonomous Mobile Robots*. Cambridge, MA: MIT Press.
- Simsek, Meryem, Adnan Aijaz, Mischa Dohler, Joachim Sachs, and Gerhard Fettweis. 2016. "5G-Enabled Tactile Internet." *IEEE Journal on Selected Areas in Communications* 34 (3): 460–473.
- Sitthi-Amorn, Pitchaya, Javier E. Ramos, Yuwang Wangy, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. 2015. "MultiFab: A Machine Vision Assisted Platform for Multi-material 3D Printing." *ACM Transactions on Graphics* 34 (4): 1–11.
- Skylar-Scott, Mark A., Jochen Mueller, Claas W. Visser, and Jennifer A. Lewis. 2019. "Voxelated Soft Matter via Multimaterial Multinozzle 3D Printing." *Nature* 575 (7782): 330–335.
- Soni, Mahesh, and Ravinder S. Dahiya. 2020. "Soft eSkin: Distributed Touch Sensing with Harmonized Energy and Computing." *Philosophical Transactions of the Royal Society A* 378 (2164): 20190156.
- Tomo, Tito Pradhono, Massimo Regoli, Alexander Schmitz, Lorenzo Natale, Harris Kristanto, Sophon Somlor, Lorenzo Jamone, Giorgio Metta, and Shigeki Sugano. 2018. "A New Silicone Structure for Uskin—a Soft, Distributed, Digital 3-Axis Skin Sensor and Its Integration on the Humanoid Robot iCub." *IEEE Robotics and Automation Letters* 3 (3): 2584–2591.
- Truby, Ryan L., Robert K. Katzschmann, Jennifer A. Lewis, and Daniela Rus. 2019. "Soft Robotic Fingers with Embedded Ionogel Sensors and Discrete Actuation Modes for Somatosensitive Manipulation." In *2019 2nd IEEE International Conference on Soft Robotics*, 322–329. New York: IEEE.
- Viteckova, Slavka, Patrik Kutilek, and Marcel Jirina. 2013. "Wearable Lower Limb Robotics: A Review." *Bio-cybernetics and Biomedical Engineering* 33 (2): 96–105.
- Wang, Xuewen, Yang Gu, Zuoping Xiong, Zheng Cui, and Ting Zhang. 2014. "Silk-Molded Flexible, Ultrasensitive, and Highly Stable Electronic Skin for Monitoring Human Physiological Signals." *Advanced Materials* 26 (9): 1336–1342.
- Weiner, Pascal, Caterina Neef, Yoshihisa Shibata, Yoshihiko Nakamura, and Tamim Asfour. 2020. "An Embedded, Multi-modal Sensor System for Scalable Robotic and Prosthetic Hand Fingers." *Sensors* 20 (1): 101.
- Yeo, Joo Chuan, Hong Kai Yap, Wang Xi, Zhiping Wang, Chen-Hua Yeow, and Chwee Teck Lim. 2016. "Flexible and Stretchable Strain Sensing Actuator for Wearable Soft Robotic Applications." *Advanced Materials Technologies* 1 (3): 1600018.
- Yogeswaran, Nivasan, Wenting Dang, William T. Navaraj, Dhayalan Shakthivel, Saleem Khan, Emre Ozan Polat, Shoubhik Gupta, et al. 2015. "New Materials and Advances in Making Electronic Skin for Interactive Robots." *Advanced Robotics* 29 (21): 1359–1373.
- Yogeswaran, Nivasan, William T. Navaraj, S. Gupta, F. Liu, V. Vinciguerra, L. Lorenzelli, and R. Dahiya. 2018. "Piezoelectric Graphene Field Effect Transistor Pressure Sensors for Tactile Sensing." *Applied Physics Letters* 113 (1): 014102.
- Yousef, Hanna, Mehdi Boukallel, and Kaspar Althoefer. 2011. "Tactile Sensing for Dexterous In-Hand Manipulation in Robotics—a Review." *Sensors and Actuators A: Physical* 167 (2): 171–187.
- Zárate, Juan José, and Herbert Shea. 2016. "Using Pot-Magnets to Enable Stable and Scalable Electromagnetic Tactile Displays." *IEEE Transactions on Haptics* 10 (1): 106–112.

9 Machine Learning for Cognitive Robotics

Tetsuya Ogata, Kuniyuki Takahashi, Tatsuro Yamada, Shingo Murata,
and Kazuma Sasaki

9.1 Introduction

In recent years, the technology of deep learning has been confirmed to be effective in various fields, such as image recognition, speech recognition, and language processing, and various applied methods have been proposed (LeCun et al. 2015). Deep learning generally refers to hierarchical neural network models of multiple layers with large dimensional inputs. One of the important characteristics of this approach is that the sensory features that human experts should typically design and select based on their knowledge and experience—for example, for a computer vision algorithm—can be self-organized through the learning process. This enables the training of deep-learning models, as long as the teaching labels are given to the target data. Data sets with high-dimensional signals can be used for training. This property enables deep learning to handle various types of data, such as images, sounds, and languages, differently from the way these problems have been treated in other research areas. The performance of deep-learning models is close to that of conventional methods and can in some modalities achieve performance superior to human abilities.

There are several methods of deep learning. One of the representative ones is with the use of autoencoders. An autoencoder is a model to learn so that input and output are the same. Once input data are provided, it is classified as “unsupervised learning” because it is simply learned to reproduce it. For example, an image input of one thousand dimensions is compressed to tens of dimensions in the middle layer, and then it is decompressed to restore the original image. Here, the low-dimensional representation in the middle layer could be used for image recognition by relearning (fine-tuning).

Convolutional neural networks (CNNs) are the current driving force of deep learning. In a multilayer network, the connections between layers are usually connected with full (dense) connection patterns. In CNNs, however, the convolution layer and the pooling layer have sparser connectivity with repeated and shared parameters, and a dense connection layer is typically added at the end. In the case of image recognition, the change of position does not affect the recognition result thanks to the convolution and pooling structure. Rather, it is important to capture a subset of features. Therefore, a CNN has small neural networks (kernels) that take only certain areas of the input image. For example, the values of three-by-three pixels are multiplied by the weights and compressed into a single value. This operation is called convolution. The kernel can be designed in many different sizes

and shapes. Since the kernel reacts to certain features in the image, it slides over the entire image area to produce a compressed representation of the image. And the pooling layer compresses the image size to save memory. With repeated convolution and pooling, the final feature representation is acquired, and finally, the recognition result is output through full connection.

There is also a set of deep networks based on recurrent neural networks (RNNs). An RNN is a neural network that does not directly connect inputs to outputs in a feedforward way, as it also has feedback connections. Even if the inputs are in a similar state, the output can change according to the internal neural condition. In RNNs, not only the connection weights but also the internal states are trained to improve prediction accuracy. RNNs have an advantage for time-series learning. However, it can be difficult to obtain good performance with an RNN because it has the same problem as deep learning, gradient vanishment. Errors are eliminated because it is difficult to propagate output errors to past steps if the learning sequences are long. To solve this problem, a new type of RNN has been developed that has multiple types of neurons. Some neurons retain their internal state in the long term (slow neurons). Some neurons change their internal state in the short term (fast neurons). These are called multitimescale neurons. As a result, fast neurons learn the short time series of input value, and long-term neurons learn the sequence of these short time series. A multitimescale RNN (MTRNN; Yamashita and Tani 2008) uses continuous neurons, with internal states represented by continuous values. By adjusting the time constant of the neuron change, it is possible to create fast and slow neurons. Another commonly used type of deep RNN is the long short-term memory (LSTM; Hochreiter and Schmidhuber 1997). In addition to the weight of the current input, the LSTM neuron learns whether to accept it (Input Gate), whether to output it (Output Gate), whether to keep the current state (Forget Gate), and other various outputs used by the error back-propagation method. LSTM models now perform well, especially in natural language processing.

Various deep-learning models and applications can be used for different modalities, such as vision, audio, and tactile modalities, in cognitive robotics. Furthermore, due to the fact that various modalities can be handled in a similar framework, these can lead to the *multimodal* applications of deep learning. In particular, a robot working in the real world is a typical multimodal system with cameras, microphones, distance sensors, tactile sensors, and actuators.

This chapter provides an overview of the research that focuses mainly on the applications of deep learning for robotics. In subsequent chapters focusing on specific cognitive robotics capabilities, more examples of deep-learning models will be discussed. The first part of this chapter contains three subsections concerning the learning of visual, tactile, and language modalities and skills. The subsequent sections focus on behavior learning related to imitation learning and on reinforcement-learning approaches. The final section discusses the possibilities of deep learning and its future prospects.

9.2 Deep-Learning Model for Modality Application

9.2.1 Robot Vision

The most natural application of deep-learning technology is in the research field of robot vision. For example, Lenz, Lee, and Saxena (2015) proposed a method to output the posi-

tion and direction (four dimensions) of a hand to grasp from a distance an image of an object. Using a CNN, Yang, Li, et al. (2015) identified forty-eight kinds of objects and six types of grasping directly from a YouTube video of a human cooking and applied them to the motion of a robot.

Redmon and Angelova (2015) also used CNNs to predict the grasping position of an object from a three-dimensional RGB-D image consisting of color (RGB) and depth (D) data. Concretely, for an RGB image of 224×224 pixels, a grasping position vector of an object is labeled by human. The grip position vector has six dimensions, including the rectangular shape of the center coordinate, the rotation angle, and the grip position of the vector (figure 9.1).

The success rate is calculated using two conditions: 1) the rotation angle error is within 30° , and 2) the overlapping area ($A \cap B$) with respect to the total area ($A \cup B$) is over 25 percent. However, these criteria do not evaluate the actual motions of the robot. The success rate of grasping using a real robot is not always high.

What is important here is that information regarding the object grasping cannot be obtained from just the image of the object. The learning process should reflect the hardware (body) of the robot and the effects of the possible motion. For example, although the grip position vector shown in figure 9.1 is a feature for a gripper, there is no guarantee that it is a sufficient and optimum feature quantity for general gripper mechanisms. When extracting a region for grasping an object, a robot should consider the physical features of the target object, such as the weight, center-of-mass, surface friction, shape change, and so on. Even if the same hand is used, grasping should be changed in various ways depending on the hand size, the payload, the direction of the approach (trajectory), and more. That is, the learning process should include not only the image of the object but also the motion generated by the robot hardware.

9.2.2 Tactile Learning

Learning the tactile sense is important for robots to allow them to obtain physical information while interacting with environments. This can be useful for operations such as




Figure 9.1
The grip position vector.

walking, physical contact with people, and object manipulation. The improved availability of tactile sensors has enabled research in this field to flourish (see chapter 8). Prior to the use of learning-based approaches, tactile sensor data were only used with handcrafted features (Yang, Sun, et al. 2016) or to trigger specific actions (Yamaguchi and Atkeson 2016). However, such methods may not scale well as tactile-sensing technology advances—for example, when a higher resolution and a larger amount of data are necessary, or as task complexity increases. By using learning-based approaches, in particular deep learning, it is now possible to handle tasks such as image recognition and natural language processing, which involve high-dimensional data and were previously difficult to process. Moreover, deep-learning approaches have recently been applied to tactile sensing, such as object recognition (Schmitz et al. 2014; Baishya and Bäuml 2016), tactile properties recognition (Gao et al. 2016; Yuan, Wang, et al. 2017), and grasping (Calandra et al. 2018).

In recent years, within research involving tactile sensors, object manipulation using robotic hands has been gaining attention since manipulation is one of the fundamental functions for a robot to perform various tasks such as tidying up, cooking, and folding clothes. In this chapter, the recent development of tactile learning and the following four categories of the object manipulation process are described: 1) object recognition, 2) grasping, 3) in-hand object pose estimation, and 4) in-hand object manipulation.

Types of tactile sensors

Many different tactile sensors have been developed to improve manipulation in robotic hands (Dahiya et al. 2013; see also chapter 8 for a detailed analysis). The majority of these sensors, however, belong to one of the following three categories:

- Multitouch sensors that can only sense force information along one axis—namely, perpendicular to the surface of the sensor. These types of sensors are known as pressure sensors (Ohmura, Kuniyoshi, and Nagakubo 2006; Iwata and Sugano 2009; Mittendorf and Cheng 2011; Fishel and Loeb 2012).
- Three-axis sensors that can sense both shear and pressure forces but are only single touch (Paulino et al. 2017).
- Three-axis sensors for both shear and pressure forces that are multitouch (Tomo et al. 2018; Yamaguchi and Atkeson 2016; Yuan, Dong, and Adelson 2017).

At the time of writing, there are only three sensors of the last type: uSkin (Tomo et al. 2018), Finger Vision (Yamaguchi and Atkeson 2016), and GelSight (Johnson and Adelson 2009; Dong, Yuan, and Adelson 2017; Yuan, Dong, and Adelson 2017). The uSkin measures the deformation of silicon during contact by monitoring changes in the magnetic fields of magnets in silicon. The sensor is able to measure both pressure as well as shear force per sensor unit for multiple contact points.

Instead of a magnet, the Finger Vision is a vision-based tactile sensor, meaning that it uses a camera to capture and measure the deformation of its attached marker during contact with a surface. In addition to contact sensing, it can also function as a proximity sensor since the Finger Vision uses transparent silicon.

The GelSight can be manufactured by covering the silicon surface of the Finger Vision with another layer of silicon that contains aluminum powder. The aluminum powder highlights the deformation of the silicon layer more clearly and hence allows for richer informa-

tion during sensing. The GelSight can be duplicated easily and is suitable for deep learning because it also uses a camera, so existing image-processing techniques can be employed to process the data. Therefore, the GelSight has become increasingly popular in research (Calandra et al. 2018; Tian et al. 2019; Zhang et al. 2020; Anzai and Takahashi 2020).

Object recognition

One of the main approaches to recognizing the type of object in a robotic hand (Schmitz et al. 2014), its materials (Baishya and Bäuml 2016; Yuan, Zhu, et al. 2017), and its properties (Gao et al. 2016), using touch and image information, is its classification through supervised learning using manually designed labels. Baishya and Bäuml (2016) and Yuan, Zhu, et al. (2017) estimated the hardness of an object as a continuous value using a tactile sensor through supervised learning. In these approaches, however, the results of class labels and their degrees depend completely on the manner in which these class labels are designed. On the other hand, one of the approaches without manually specified labels represents tactile properties in a continuous space using an unsupervised-learning approach (Takahashi and Tan 2019).

Grasping

A different use case is shown in Calandra et al. (2018), in which they utilized deep reinforcement learning and combined input data acquired from a tactile sensor with images to grasp objects using a parallel gripper, which improved their success rate in grasping experiments compared to only vision. Wu et al. (2019) showed similar results using a multifinger hand. By using a tactile sensor, the stability of a grasp can be evaluated and improved upon regrasping (Calandra et al. 2018; Wu et al. 2019; Hogan et al. 2018).

In-hand object pose estimation

In order to realize the target object pose, it is necessary to be able to estimate the current object posture. Object pose estimation is a well-studied problem in computer vision. Many researchers have been developing methods using depth data (point cloud) or RGB-D data (Choi and Christensen 2012; Aldoma et al. 2012; Choi et al. 2012). Classical approaches with depth data are mainly based on point cloud matching methods, such as iterative closest point (ICP; Rusinkiewicz and Levoy 2001). Since this method requires three-dimensional (3D) models of objects, unknown objects cannot be handled. In the state-of-the-art research in pose estimation, methods that do not require 3D models have been studied using deep learning (Schwarz, Schulz, and Behnke 2015; Hodaň et al. 2018; Hu et al. 2019).

These methods, however, are challenging to apply to in-hand manipulation because of occlusion by the hand in the image or depth data. Since tactile sensors can observe the contact state despite a visual occlusion, they are suitable for overcoming this challenge. Some research has performed object pose estimation with tactile sensors by means of a model-based approach using a 3D model (Bimbo et al. 2016) and without using a 3D model (Anzai and Takahashi 2020).

To overcome challenges such as occlusions or lack of sufficient information, one can use multiple sensors to try to obtain an improved perception of the environment or situation. In this case it is of great importance to know which modals can be trusted in a given situation—in other words, how reliable a given sensor modal is. For example, if a vision sensor is impaired, one should give its data less importance than other sensor modals. It is difficult, however, to determine sensor modal reliability through rule-based methods.

Anzai and Takahashi (2020) proposed a network that can autonomously determine the reliability of each modal.

In-hand object manipulation

To manipulate a grasped object to a target posture is one of the most challenging tasks. Analytical approaches exist, but they come with limitations, such as the known object model and the rigid object (Han et al. 1997; Han and Trinkle 1998). In learning-based approaches, manipulation is performed by predicting the state of the tactile sensor for the motion of a robot's end effector (Tian et al. 2019; Li et al. 2014; Funabashi et al. 2018). Since object manipulation with a multifingered hand is still challenging, most of these studies are simple tasks and take place in experimental settings, with a few exceptions (e.g., Falco et al. 2018).

9.2.3 Learning of Language Grounding in Robot Behavior

Natural language is the most powerful tool for expressing our requests to other agents. Service robots must be able to understand natural language to flexibly respond to human requirements or to effectively work together with humans. However, to arbitrarily design mapping between language, which is a discrete system, and the referents in the real world, which is a continuous and dynamical system, is notoriously difficult, as stated in the symbol grounding problem (Harnad 1990). The meanings of linguistic expressions also greatly depend on the current context that an agent is situated in. For instance, to respond to the instruction “grasp the red ball,” a robot is required to generate different trajectories of joint angles in accordance with the position of the red ball. Unlike most situations in industrial factories, our living environment is highly changeable and open ended; new situations almost always differ from the previous ones. It is almost impossible to make explicit rules that can handle all possible situations in a top-down manner.

Many attempts have been made to get robots to learn grounding relationships from their own experiences in a bottom-up manner. Here we review existing studies that consider the learning of grounding relationships between language and behavior in robots. In particular, we discuss the two main approaches to language grounding: probabilistic modeling and neural networks. See also chapter 20 for more details on deep-learning approaches to robot language models.

Probabilistic modeling

One way to model the relationships between language and other modalities is to model them as probabilistic relationships. For example, Inamura et al. (2004) utilized hidden Markov models (HMMs) to recognize and generate human motions. In their framework, protosymbols, which represent a specific motion pattern, emerged in the learning process. Nishihara, Nakamura, and Nagai (2017) utilized a multimodal latent Dirichlet model (MLDA) for a robot to learn object concepts that connected multimodal information consisting of co-occurring word, auditory, visual, and tactile data. Tellex et al. (2011) proposed a framework called generalized grounding graphs, which dynamically instantiated a graphic model depending on the semantic structure of linguistic commands, and they then inferred appropriate plans for navigation and manipulation in the graph.

One advantage of probabilistic models is their high intelligibility. In the case of graphic models, each node in the graph is designed as a meaningful element. Therefore, it is easy

to understand what kind of inference is performed by the model. However, a probabilistic model that has the capability of dealing with long-term dependencies sufficiently has not yet been developed.

Neural networks

On the other hand, methods that model language grounding deterministically also exist. One popular method is neural networks, such as with RNNs. Sugita and Tani (2005) proposed a trainable architecture that consisted of two neural networks—one of which was for language and the other, robot behavior—with a small number of shared nodes called parametric bias (PB). The model learned to embed the relationships between language and behavior in topological organization in the PB space. Ogata et al. (2007) employed a similar architecture to learn the bidirectional mapping between language and robot behavior. Heinrich and Wermter (2014) proposed a model that connected three RNNs. Each RNN was specialized for vision, proprioception, and language, respectively, but they were connected to each other. After learning, the model could generate sentences that described robot motions as a sequence of characters. Stramandinoli, Marocco, and Cangelosi (2017) utilized a Jordan-type RNN (Jordan 1997) to ground abstract words (e.g., use and make) in robots' sensorimotor experiences. The abstract words were learned by recalling the meanings of previously learned basic words and combining them.

An advantage of neural networks is that by introducing recurrent connections and some gating mechanism, such as LSTM (Hochreiter and Schmidhuber 1997), they can achieve a much higher performance in learning temporal structure with long-term dependency without a priori knowledge. One disadvantage of neural networks is that it is difficult to understand their behavior since their representations in hidden layers are in a distributed form. Recently, some studies have proposed methods to visualize the internal behavior of neural networks (Bach et al. 2015; Smilkov et al. 2017) and to make their representations more intelligible (Chen et al. 2016; Xu et al. 2015). The following introduces a recent study that proposed an RNN-based framework to ground language in robot behavior.

Yamada, Matsunaga, and Ogata (2018) attempted to bidirectionally convert language and robot behavior by utilizing two coupled recurrent autoencoders (RAEs; figure 9.2): one RAE coped with language, and the other dealt with behavior.

Each RAE consists of an encoder RNN and a decoder RNN. The encoder RNN compresses a time series (a sentence or a behavioral sequence; x_1, x_2, \dots, x_T) into a fixed-dimensional feature vector z :

$$z = \text{EncoderRNN}(x_1, x_2, \dots, x_T)$$

The decoder RNN produces a sequence by recursively decoding the feature vector:

$$(y_1, y_2, \dots, y_T) = \text{DecoderRNN}(z)$$

The RAE is trained to reconstruct the original sequence through the feature vector—namely, identity function. The loss function is as follows:

$$L = \frac{1}{T} \sum_{t=1}^T \psi(x_t, y_t).$$




Figure 9.2

Two coupled RAEs to bidirectionally convert language and robot behavior. *Source:* Adapted from Yamada, Matsunaga, and Ogata 2018.

The detail of loss function ψ at each time step depends on the modality. In the learning process, the language RAE and the behavior RAE are optimized to extract the important features of time series data in each modality.

In addition, the whole system is trained in such a way that the feature vectors of co-occurring language and robot behavior get closer to each other, and the feature vectors of unpaired language and behavior grow more distant from each other. With this constraint, this coupled RAE system is able to bidirectionally convert language and behavior through the latent feature space. Producing a behavior sequence in response to a sentence is realized by using the encoder of the language RAE to encode the sentence and the decoder of the behavior RAE to expand the feature vector. In contrast, producing a sentence description of a robot behavior is realized by having the encoder of the behavior RAE encode a behavioral sequence and having the decoder of the language RAE expand the feature vector.

Figure 9.3 shows the latent feature spaces organized by learning in this robot experiment. Each point corresponds to a sentence in the left panel and to a behavioral sequence in the right panel. It can be seen that the behavioral sequences were actually bound with their paired sentences. Here, it is worth noting that because the behavior RAE also receives vision input, the model could respond to the same sentence by producing different joint-angle trajectories depending on the current contexts.

9.3 Imitation Learning (Predictive Learning)

Imitation learning, also referred to as learning from demonstration (LfD) or programming by demonstration (PbD), is a learning-based approach that enables robots to acquire skills (or infer policies) for action generation from a set of expert demonstrations representing the robots' sensorimotor experiences. Imitation learning is mostly performed by a scheme of predictive learning in which robots are required to learn to predict the (sensory-)motor state at the next time step from the sensory(-motor) state at the current time step. This is a more data efficient approach in comparison to the reinforcement learning to be introduced in a forthcoming section. Imitation learning is a particularly useful approach when the use of a reinforcement-learning algorithm is unrealistic due to the difficulty in design-

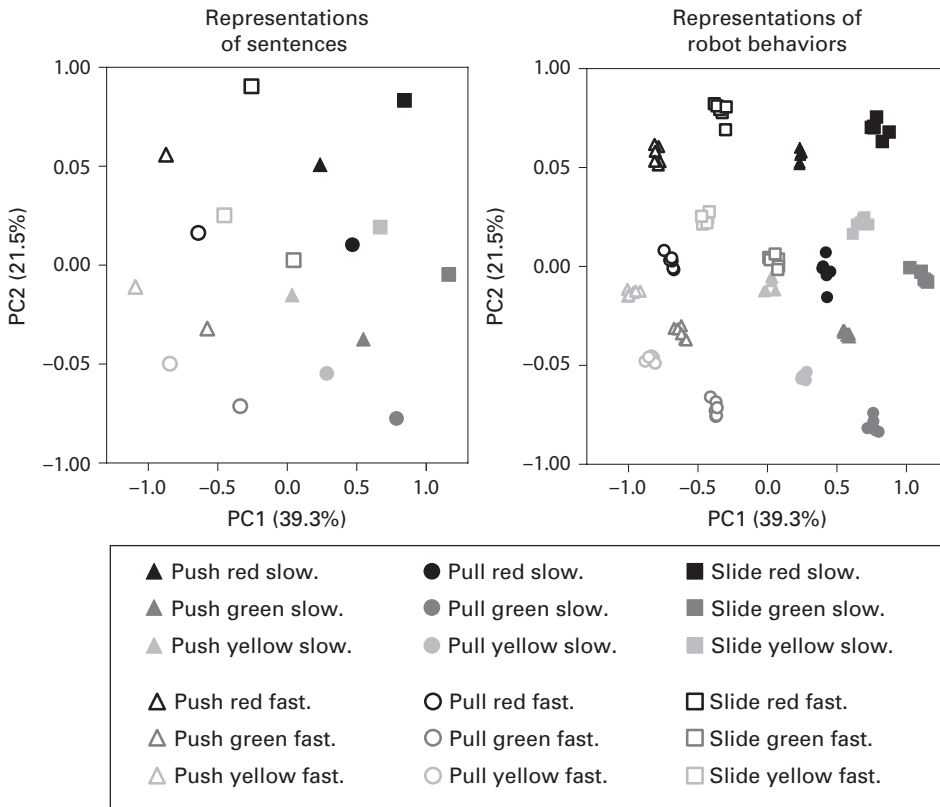


Figure 9.3

Latent representations of language and robot behavior by the coupled RAEs. *Source:* Adapted from Yamada, Matsunaga, and Ogata 2018.

ing a reward function and in performing a massive amount of exploration (with real robots). In the context of (cognitive) robotics and robot learning, imitation learning includes the following two cases: 1) learning from sensorimotor experiences and 2) learning from sensorimotor experiences by observing another agent's demonstrations. In both cases, it is necessary to provide demonstrations about robot performance during the learning process via kinesthetic teaching or teleoperation by a human demonstrator. The difference between them is whether or not the sensory (mainly visual) experiences include demonstrations about the performance of another agent, typically a human. Namely, in the second case robots are required not only to learn to generate their own actions but also to map an observed other's actions to their own by inferring what to perform and how to perform. This is much closer to the original meaning of imitation by humans and animals in the context of cognitive science (Meltzoff and Moore 1977).

There are several machine-learning approaches for performing imitation learning, such as neural networks (e.g., CNNs and RNNs); probabilistic models such as the combination of a Gaussian mixture model and Gaussian mixture regression (e.g., Calinon, Guenter, and Billard 2007); hidden Markov models (e.g., Inamura et al. 2004); and dynamical systems (e.g., dynamic movement primitives in Ijspeert, Nakanishi, and Schaal [2002] and Ijspeert

et al. [2013]). In this section, we focus particularly on neural network-based approaches (refer to review papers for other approaches, such as Argall et al. [2009] and Billard et al. [2008]). In what follows, several studies of the above two cases of imitation learning are examined. In addition, their extensions with deep-learning approaches, such as the use of deep autoencoders for visual feature extraction from raw images and LSTM for learning long-term dependencies, are introduced. Finally, related advanced topics, including one-shot imitation learning and self-supervised learning from play data, are also briefly discussed.

9.3.1 Imitation Learning from Own Sensorimotor Experiences

Ito et al. (2006) studied the learning of primitive actions for object manipulation by using an RNN with parametric bias (RNNPB). In their experiment, the sensorimotor experiences of a small humanoid robot QRIO for ball handling were first collected via kinesthetic teaching. There were two different primitive actions for ball handling, including: 1) rolling a ball from the left to right sides and vice versa (referred to as ball-rolling action hereafter) and 2) lifting the ball and letting it fall to the ground (referred to as ball-lifting action hereafter). Sensorimotor experiences consisted of time-series data items (or trajectories) of visual information represented as ball position and action information represented as joint angles of both arms. The robot with an RNNPB was required to learn to predict the visuomotor state at the next time step given the state at the current time step. Through this learning process, the various primitive actions were represented by the difference in optimized PB vectors. Namely, once a PB vector corresponding to the ball-rolling action is set into the network, the robot generates the ball-rolling action, and once the other vector corresponding to the ball-lifting action is set, the robot generates the ball-lifting action. This means that different primitive actions were acquired as multiple limit cycle attractors in the RNNPB.

One of the important points of this experiment is that the PB vector during action generation after the learning phase was also optimized online in the direction of minimizing prediction errors computed during a time window of immediate past time steps. This iterative optimization of the PB vector enabled the robot to adapt to unexpected situational changes. For example, consider a situation in which the PB vector for the ball-rolling action is set, and the robot is generating the corresponding action. Then, an experimenter suddenly disturbs the ball movement between the left and right sides, and the ball movement stops at the center front of the robot. Before the disturbance, the robot was predicting that the ball would be moving between the left and right sides as a consequence of its own action generation. However, due to the disturbance that stopped the ball movement, the robot feels a discrepancy between the anticipated and actual situations or prediction errors. The only solution to minimize these errors is to switch the originally set PB vector to the other one that generates the ball-lifting action. This switching of the PB vector enables the robot to minimize the generated prediction errors and to perform stable action generation again. The important point of this phenomenon is that the robot had never learned to switch between the different primitive actions. Thanks to the simple computational principle of the so-called prediction error minimization (Nagai 2019), the robot realized adaptive action generation. This is closely related to the active inference scheme based on the free energy principle (Friston et al. 2010).

Chen, Murata, et al. (2016) extended the framework to an interaction between two NAO robots. In their experiment, each robot with an RNNPB first learned a set of primitive actions for ball manipulation with a human experimenter. The learned primitive actions were dependent on the ball movement such that when the ball was heading toward the right side of a robot, the robot was required to hit the ball with its right hand. After the learning phase, the robots faced each other and were required to perform a ball-play interaction. Because the experiment was performed in the real world, with some fluctuations such as the friction between the ball and a table, sometimes the ball dynamics suddenly changed in an unpredictable manner. In such a situation, prediction errors arose in both the robots, and these errors triggered the PB vector of each robot, optimizing it to fit the current situation. This dual optimization of the PB vector of each robot enabled spontaneous action switches without any training.

In the former examples using an RNNPB, the switch between primitive actions was triggered by environmental changes. Next, we consider how such switching can be intentionally generated by learning action sequences consisting of combinations of primitive actions. Yamashita and Tani (2008) and Nishimoto and Tani (2009) tackled this issue by using the MTRNN introduced above. In a manner similar to the RNNPB experiments introduced earlier, they first collected visuomotor experiences of the QRIO robot via kinesthetic teaching. The recorded sequences were more complex than the first study above. For example, in one sequence the robot reached for an object from a home position and then moved the object up and down three times before finally moving it back to the home position. Specifically, each sequence contained multiple primitive actions such as reaching for and moving the object, and the robot was required to switch or repeat such actions. The robot with an MTRNN performed predictive learning of these complex and longer visuomotor experiences by utilizing the sensitivity of the initial conditions of the slow dynamics layer of the MTRNN. After the learning phase, the robot succeeded in generating the learned action sequences. Analysis of the fast and slow dynamics layers revealed that primitive actions were represented in the fast dynamics layer, and the combinations of these primitives (sequence information) were represented in the slow dynamics layer thanks to the self-organized functional hierarchy.

Namikawa, Nishimoto, and Tani (2011) extended this experimental setup and considered how probabilistic transitions among primitive actions could be learned. In the same manner as the former cases, they first recorded visuomotor experiences for an object manipulation in which the QRIO robot moved an object from center to left, from left to center, from center to right, and so on via kinesthetic teaching. These transition patterns were determined probabilistically, and they investigated whether such sequences with probabilistic transitions could be learned by a deterministic MTRNN. The robot after the learning phase reconstructed a demonstrated visuomotor sequence from the beginning by setting an optimized initial state of the slow dynamics layer, but the sequence gradually changed from the learned one. The analysis of the generated action sequences demonstrated that the transition probabilities were still preserved in newly generated sequences. The analysis of each layer of the MTRNN revealed that in the same way as in the former studies (Yamashita and Tani 2008; Nishimoto and Tani 2009), different types of information were stored in each layer. One more interesting phenomenon is that only the slow

dynamics layer exhibited chaotic dynamics with a positive Lyapunov exponent, which led to the reconstruction of the probabilistic transitions by deterministic neural dynamics.

In the experiments conducted before the deep-learning era, such as that just described, the experimental setup was simplified so that, for example, the visual information was just the object position. Here, some scaled-up experiments are introduced that deal with high-dimensional raw visual images by using deep-learning approaches such as a deep (convolutional) autoencoder.

Noda et al. (2014) conducted a study on the integrative learning of multimodal information such as vision, auditory, and motor data using a combination of deep autoencoders for feature extraction and temporal processing. As in the previous studies, they first collected sensorimotor experiences of the NAO robot via kinesthetic teaching. Then, low-dimensional features of high-dimensional raw visual images and auditory information were extracted by using the respective deep autoencoders. The extracted visual and auditory features were concatenated with joint angle information. They used another deep autoencoder called a time-delay neural network (TDNN) that received a time window of the multimodal information and outputs its reconstruction. By using this framework, they realized action generation by prediction and retrieval, such as visual retrieval from auditory and joint angle information using high-dimensional sensorimotor states.

Yang et al. (2017) extended this framework to the human-size industrial robot Nextage and performed a towel-folding task. It is known that towel handling is a challenging task in robotics because modeling a deformable object is difficult. They recorded visuomotor experiences via teleoperation using a 3D mouse. In their experiment, the normal autoencoder for visual feature extraction was replaced with a deep convolutional autoencoder (ConvAE). They realized repeatable towel folding with a high success rate after the learning phase. Kase and colleagues replaced the TDNN used in the above two experiments with RNN-based architectures, an MTRNN (Kase et al. 2018) and an LSTM (Kase et al. 2019). These replacements realized much longer and complex task executions such as put-in-the-box and skewering thanks to their characteristics of functional hierarchy and long short-term memories.

9.3.2 Imitation Learning from Observing Another Agent's Demonstrations

When learning from observing another agent's action generation, robots need to infer what to perform and how to perform. Arie et al. (2012) considered this issue by using an MTRNN. In their experiment, a small humanoid robot, HOAP-3, learned a set of visuomotor sequences consisting of multiple primitive actions. For example, in one sequence the robot first reached for an object from a home position, then moved the object right, then knocked the object over, and finally moved the object back to the home position. Note that the robot learned not only its own action generation but also how to map an observed action of human performance to its performance. There were four primitive actions, including R (moving the object to the right), L (moving the object to the left), K (knocking over the object), and U (moving the object upward). The robot first learned three different types of visuomotor sequences (RK, UK, and UL) produced by itself and the experimenter. After these sequences, the robot was subjected to the demonstration of only the human's performance for the RL sequence. The robot was evaluated on whether it could generate its

own action for the RL sequence, which had not been learned, by mapping the observed demonstration of human performance to its own performance.

The slow dynamics layer of the MTRNN had two special neural units whose initial conditions were optimized to be the same values when the demonstrations were the same patterns, regardless of the generation of robot performance and the observation of human performance. The other two units in the slow dynamics layer served as a PB vector that discriminated the self-mode (generation of robot performance) and the other-mode (observation of human performance) by assigning a particular value for each (one for the self-mode and minus one for the other-mode). In the evaluation after the additional learning phase, an action-specific initial state for the demonstration of human performance for the RL sequence was set, and the demonstrator-specific PB vector was switched to the self-mode. This enabled the robot to generate the unlearned combinatory actions for the RL sequence.

Nakajo et al. (2015) considered another important topic concerning the acquisition of viewpoint representation. Humans can understand what action is demonstrated by another regardless of a difference in viewpoint. Acquiring such an ability is useful for robots because the demonstration of human performance can be provided from any direction. However, this is not straightforward for robots because the visual information from the demonstration of human performance from different viewpoints is distinct. They used an MTRNN for learning the demonstrations of object manipulation for both the robot and human performances. In their experiment, a human demonstrator performed actions from multiple viewpoints. They provided constraints on the initial state optimization by introducing a subnetwork for representing viewpoints. Their analysis of the initial state space of the subnetwork revealed that the positional relationship of the viewpoints was self-organized in the space. In their experiment, although the structured representation of viewpoints was self-organized, how to map the demonstration of human performance provided from multiple viewpoints to the same robot performance remained an issue.

To tackle this issue, Nakajo et al. (2018) extended the experiment by introducing a sequence-to-sequence (seq2seq) deep-learning approach that has been widely used, especially in machine translation (Sutskever, Vinyals, and Le 2014). The seq2seq framework consists of an RNN-based encoder-decoder architecture. In the machine translation, the encoder RNN receives source sentence information, such as an English sentence, sequentially and transforms it into a fixed-dimensional vector. The decoder receives this vector and transforms it to target sentence information, such as a Japanese sentence. By referring to this information processing of the seq2seq framework, they first encoded visual features of video information about the demonstration of human performance extracted by a convolutional encoder with an MTRNN. Then an achieved fixed-dimensional vector was transformed to the robot's action generation. After a learning phase, the robot was able to map the demonstration of human performance provided from an unlearned viewpoint to its own action generation. The analysis of each layer of the MTRNN shows the representation of actions, objects, and viewpoints. More specifically, after the demonstration of human performance, the fast dynamics layer represented viewpoint information, and the slow dynamics layer represented action and object information without any viewpoint information. The key point for the success of mapping from unlearned human demonstration

to robot performance is that the slow dynamics layer acquired the viewpoint-invariant representation about the actions and objects by squishing the viewpoint information, which is unnecessary for a robot's own action generation after the observation.

9.3.3 One-Shot Imitation Learning and Self-Supervised Learning

One of the new directions in imitation learning is one-shot imitation learning (Finn et al. 2017; Yu et al. 2018; Duan et al. 2017). One-shot imitation learning means that robots are required to learn a new task from only a single demonstration of the robot's or human's performance for the given task. As an example, Finn et al. (2017) combined a metalearning algorithm called model-agnostic meta-learning (MAML; Finn, Abbeel, and Levine 2017) and imitation learning. The MAML enables neural networks to learn a new task from only a few training data. More specifically, the MAML assumes various tasks, and it samples some tasks from which it also samples training and validation data items (at least one item for each). During a meta-learning phase, first the training loss for each task is computed by using initial model parameters and the sampled training data item. By using the computed training loss for each task, the initial model parameters are (tentatively) adapted for each task by gradient descent. Then the validation loss for each task is computed by using the corresponding adapted parameters and the sampled validation data item. Finally, the initial model parameters are optimized to minimize the sum of the validation losses by gradient descent. This means the metalearning algorithm tries to discover generalized initial parameters that can be easily adapted for any task. During a subsequent meta-testing phase, only a single training data item from a new task kept separate from tasks for the meta-learning phase is given, and the generalized initial parameters can be quickly adapted to the task.

In their experiment using a robot PR2, they first collected demonstrations of robot performance for various tasks of object placing via teleoperation. The collected demonstrations consisted of raw visual images from a camera mounted on the robot and action information. The meta-learning was conducted by using these demonstrations to learn how to infer a policy for a new task from only a single demonstration of robot performance. Then, in the meta-testing phase, the robot learned a new task from a single demonstration provided via teleoperation by a human. This is effective for learning a new task quickly; however, the problem is that the framework needs a demonstration of robot performance, and providing a single demonstration of human performance is more straightforward. To tackle this issue, Yu et al. (2018) extended the framework by introducing domain-adaptive meta-learning (DAML). This enables robots to learn how to infer a policy for a new task from only a single demonstration of human performance. They evaluated this extended framework with both the PR2 and Sawyer robots. As expected, these robots could learn a new task from a single demonstration of human performance and could also learn a new task even when the demonstration was performed in different viewpoints and background environmental situations.

Another new direction is self-supervised learning (Nair et al. 2017; Pathak et al. 2018; Lynch et al. 2019). In all the experiments explained above, the demonstrations by human experts were provided for performing specific tasks. As an alternative approach, Lynch et al. (2019) proposed a new paradigm of learning from play (LFP), in which robots acquire

various skills for object manipulation only from play data given by teleoperators and realize goal-directed tasks after a learning phase. In their experiment, human operators first teleoperated a robot in a simulation environment. In the environment, multiple objects for manipulation sat on a desk equipped with a drawer and a shelf with buttons that turned on lights. The operators were asked to freely explore the environment by operating the robot, and visuomotor experiences during this free exploration were collected. The important point is that the curiosity and intrinsic motivation of the operators enabled the acquisition of various types of complex and interactive actions with both manipulative and nonmanipulative objects available in the environment. The collected visuomotor experiences were learned by the play-supervised latent motor plans (Play-LMP) framework that consists of a plan proposal encoder, a plan recognition encoder, and an action decoder. During a learning phase, the first part of the visuomotor experiences was randomly sampled as a sequence. Then only the initial and final states of the sampled sequence were encoded by the plan proposal encoder, and the entire sequence was encoded by the plan recognition encoder. Both encoders generated a latent plan representation and that from the recognition encoder was provided for the action decoder. The encoders and decoder were jointly optimized to maximize action likelihood on the decoder and minimize the KL divergence between the distributions of the latent plan representations from the encoders. After the learning process, providing the current and goal states to the plan proposal encoder and sending the generated latent plan representation from this encoder to the action decoder can generate an action sequence that interpolates the current and goal states. The experimental results showed that the robots that learned from play data were more robust to perturbations in comparison to robots that learned from demonstrations for specific tasks. They also exhibited retrying-until-success behavior thanks to the diversity of the play data.

9.4 Reinforcement-Learning Robot Applications

In the previous part of this chapter, we reviewed neural network-based methods to control robots using predefined data sets of a robot's behavior. In contrast to this "off-line" method, online learning techniques collect samples of the training data set while optimizing models. We now take a look at online learning methods with the deep-learning method called "deep reinforcement learning." This approach provides a way to explore solutions that enable a robot to learn visuomotor tasks instead of a carefully designed training data set. However, it is known that reinforcement-learning methods tend to require large amounts of episode sampling because of noises of rewards or the stochastic property of interaction. In the case of robot tasks, performing many episodes with real robots is costly (e.g., time, computational costs, robot hardware reliability). In this section, we first give an overview of the reinforcement-learning problem setting. Next, we review research on robot tasks using deep reinforcement learning from the viewpoint of how to reduce the cost of episode sampling.

9.4.1 Reinforcement-Learning Problem Setting

The reinforcement-learning (RL) problem setting assumes the interaction between a controllable agent (e.g., a robot controller) and an environment (Sutton and Barto 2018; figure 9.4). For example, a controller of a picking robot can be regarded as an agent, and the environment

corresponds to the space surrounding the robot with some target objects. The agent interacts with the environment by performing an action a . Then the environment's states are altered by the action, and this returns new states and a reward signal r . The reward signal represents how well the current state transition is going, such as the achievement of the task—for example, it may be +1 when the robot successfully picks an object, 0 when the robot moves its arm toward the object, and -1 for failures. The interaction between the agent and the environment will produce a sequential tuple of state, action, and new state with reward (s, a, r, s') . Usually, the RL problem assumes this tuple is sampled from a finite Markov decision process (MDP). To infer an action from the current state is represented as a function called policy $\pi(a|s)$, and the state transition dynamics is formulated as a stochastic probability function $p(s'|s)$. The goal of RL is to find a policy that can maximize the expected sum of reward (called return) in each state of interactions. The expected return is often called “value” $v(s) = \mathbb{E}(\sum r|s)$.

Finding the best policy or explicitly computing the accurate value is intractable due to the stochastic property of MDP; thus, we need to approximate value function. RL approaches can be categorized into several types of this approximation method. One of them is to approximate value conditioned by actions, called “action value.” If we can compute an accurate action value, the agent will be able to obtain the best return by selecting an action whose action value is the highest at each time step. The action value is also difficult to compute as well as the state value, so it should be approximated by Monte Carlo methods on episode data sampled by the interactions between agent and environment. The RL Q-learning method adopts a bootstrapping method of the action value by predicting the sum of discounted future rewards. The approximation ability of the action value estimator is the key to the performance of Q-learning. Using deep-learning models as action value approximators has led to significant improvement in RL agents' abilities in video game environments, whose states are usually large-dimensional image data (Mnih et al. 2015; Vinyals et al. 2019). The other RL approach is to optimize a parameterized policy function directly. In the context of deep RL, the policy function is implemented using deep-learning models and optimized via gradient ascent toward the higher state value, called the policy gradient method. This optimization method allows actions to be in continuous space, whereas Q-learning usually allows only discrete action space. Policy gradient methods




Figure 9.4
Interaction between agent and environment.

have several variants with respect to the type of policy functions and optimization techniques used to stabilize value estimation. Another way to categorize the RL approach is to distinguish whether a learning method is explicitly modeling state transition probability $p(s'|s)$. Methods that model state transition probability are called “model based,” whereas “model-free” do not model it. The model-based approach promises lower sample complexity compared to model-free methods because we could substitute predicted future states for states given by running real interactions. When the action space is discrete and the state transition can be accurately simulated on a long time-step horizon, the heuristics of action searches, such as the Monte Carlo tree search, can be used for collecting good sample data for value estimation (Silver et al. 2018). In cases of robotic experimental settings, the state is required to have a large amount of sensory data, including camera images or poses of the robot, so other value estimation or policy optimization methods are required.

By harnessing the power of the deep-learning model’s function approximator ability, RL methods have recently been applied to large-dimensional state data and complex tasks, such as games (Mnih et al. 2015; Vinyals et al. 2019) and generative tasks (Ganin et al. 2018; Huang, Heng, and Zhou 2019), including in robotics. However, RL still requires us to collect a good deal of sample data by having the agent explore the environment, in contrast to imitating expert behavior by supervised learning. Running a lot of real robot interactions requires a huge cost in terms of the experiment and the risks of damaging the robots as they explore. Therefore, deep RL researchers have tried to make optimization methods more efficient and stable. One major research direction is to make data collection efficient, and the other is to leverage sample complexity using model-based approaches.

9.4.2 Making Data Collection Efficient

One of the ways to reduce data collection using real robots is to utilize physics simulation software. Although a simulator drastically reduces the cost of experiments, there are huge *reality gaps* due to the limited abilities of simulated environments and robots to reproduce physical world dynamics. One of the approaches to overcome the problem of the reality gap is to augment collected sample data by adding noise to simulation processes, also known as “domain randomization” (Tobin et al. 2017). For example, experiments by Andrychowicz et al. (2020) randomized the property of the robot, the physical parameters such as mass or gravity, and the visual appearance. Domain randomization is expected to improve generalization ability with regard to noise in real environment states or state transition dynamics. Instead of randomizing the state given by a simulator’s renderer, replacing state images with more realistic images faked by a generative model has also been investigated. Bousmalis et al. (2018) reported that they drastically reduce the amount of episode sampling in the real robot environment by enhancing the quality of the simulated state image using a generative adversarial network. Adding constraints to force an RL agent trained in simulated environments to behave like an agent in the real environment has also been attempted. Fang et al. (2018) incorporated the adversarial loss of classifying the source of episode data in order to transfer knowledge from an agent in simulation to one in the real environment.

RL experiments on simulators often require multiple software environments running in parallel for sampling efficiency. Conducting real robot exploration tasks in parallel could also reduce data collection time. Levine et al. (2018) built multirobot arm-picking environments

and trained an action value estimator for large collected data samples of images of cluttered objects. The action that controlled the robot arm was obtained by an evolutionary strategy whose candidates were evaluated by the action values estimated as success rates by a deep neural network.

Providing expert episode sequences helps exploration. It is also expected to reduce data collection cost. Peng and colleagues (Peng, Abbeel, et al. 2018; Peng, Kanazawa, Malik, et al. 2018) showed that human motion capture data assisted with a robot control agent's exploration in a simulator. They added a reward that encouraged simulated robots to take poses similar to a human's target poses in the original task, such as walking or performing acrobat motions. Also, the initial state at exploration was sampled from target poses to observe states that are difficult to achieve by taking random actions from the same initial state.

Incorporating reward for imitating expert sequences is related to inverse reinforcement learning, which is an RL approach for estimating reward function from expert data (Ng and Russell 2000). Finn et al. (2016) and Peng, Kanazawa, Toyer, et al. (2018) proposed the use of a generative adversarial protocol to determine the similarities between episodes by the RL agent and the expert data. In this case, the reward was given by a discriminator network trained to distinguish between the sequences from the agent's exploration and the expert. A training reward function approximator network was also expected to relieve the sparseness of the reward. Basic RL requires us to design reward functions for representing task achievements carefully. Very sparse reward distribution, such as a nonzero signal only at the end of an episode, makes exploration challenging since value estimation becomes unstable. A reward estimator by a trained machine-learning model is expected to give nonzero rewards even during episodes. Ganin et al. (2018) proposed a painting RL agent that can be trained by reward signals given by a discriminator network able to distinguish whether a picture image is drawn by the agent or by a human.

9.4.3 Reducing Data Collection by Modeling Environment Dynamics

Model-based RL methods allow policy optimization to acquire sequential data predicted from environment models, and thus they promise to reduce sample complexity in contrast to model-free algorithms. The recent success of generative deep-learning models has led to their utilization in modeling high-dimensional and complex state transitions—for example, image frame sequences. Ebert et al. (2018) proposed image sequence modeling conditioned by a robot's actions for object manipulation tasks. They collected image sequences by moving the robot's arm with random actions and training a deep convolutional network to predict future image frames. After training an image frame predictor, actions were directly optimized by a cross-entropy method, which is a derivative-free optimization method. They produced multiple predicted image sequences from their existing image obtained by a robot with action candidates. Each action candidate was then evaluated with the predicted image at the end of the time-step horizon for differences between the given goal image and the predicted image, or pixel annotation by an experimenter. A combination of future image predictions and a derivative-free algorithm were also proposed by Ha and Schmidhuber (2018). In this study, a state transition function was modeled by a stochastic neural model based on a mixture density network. They argued that the states predicted by deterministic dynamics make the policy optimization adver-

sarial. Nevertheless, nondeterministic modeling will easily lead to inaccurate state prediction due to the uncertainty of the future. Hafner et al. (2018) proposed a combination of both RL modeling methods using a recurrent state-space model (Karl et al. 2019). Instead of directly optimizing the action sequence, model-free RL methods can be used jointly with model-based RL methods. An issue when combining model-based RL with model-free optimization methods is inaccurate dynamics modeling. Kurutach et al. (2018) indicated that policy optimization tends to exploit the region of state space insufficient for achieving good performance. Buckman et al. (2018) proposed the use of an ensemble of several versions of the learned dynamics to stabilize value estimation.

9.5 Conclusion

This chapter introduced several research examples of robot applications using machine learning, especially deep learning, for tasks such as robot vision, the learning of tactile sense and motion, imitation learning, prediction learning, reinforcement learning, and language learning.

It is important to realize that robotics research showing the robot's performance only in simulation and/or in specific environments cannot lead to practical applications. One of the most critical conditions to consider is the evaluation of the robustness of the various noisy situations in the real environment.

In Japan, various manufacturers of industrial robots have already developed multiple prototypes of robot applications of imitation learning and prediction learning. The modularization of robotic systems at the hardware and software levels is progressing quickly, and big developments are expected to be realized with deep-learning technology. In general, the robotics approaches using AI deep-learning methods have the potential to significantly advance cognitive capabilities in robots.

Additional Reading and Resources

- A comprehensive book on deep-learning methods: Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press (free online copy: <https://www.deeplearningbook.org>).
- Position paper discussing the challenges and opportunities connecting robotics with deep learning: Sünderhauf, Niko, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upercroft, et al. 2018. "The Limits and Potentials of Deep Learning for Robotics." *International Journal of Robotics Research* 37 (4–5): 405–420.
- Recent volume with extensive coverage of reinforcement-learning methods: Sutton, Richard S., and Andrew G. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- OpenAI Gym tool kit for developing reinforcement-learning simulation, including with simulated robots: <https://gym.openai.com>.

References

- Aldoma, Aitor, Zoltan Csaba Marton, Federico Tombari, Walter Wohlkinger, Christian Pothast, Bernhard Zeisl, Radu Rusu, Suat Gedikli, and Markus Vincze. 2012. "Tutorial: Point Cloud Library: Three-Dimensional Object Recognition and 6 DOF Pose Estimation." *IEEE Robotics and Automation Magazine* 19 (3): 80–91. <https://doi.org/10.1109/mra.2012.2206675>.
- Andrychowicz, Open AI: Marcin, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, et al. 2020. "Learning Dexterous In-Hand Manipulation." *International Journal of Robotics Research* 39 (1): 3–20. <https://doi.org/10.1177/0278364919887447>.
- Anzai, Tomoki, and Kuniyuki Takahashi. 2020. "Deep Gated Multi-modal Learning: In-Hand Object Pose Changes Estimation Using Tactile and Image Data." In *IEEE International Conference on Intelligent Robots and Systems*. New York: IEEE.
- Argall, Brenna D., Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. "A Survey of Robot Learning from Demonstration." *Robotics and Autonomous Systems* 57 (5): 469–483. <https://doi.org/10.1016/j.robot.2008.10.024>.
- Arie, Hiroaki, Takafumi Arakaki, Shigeki Sugano, and Jun Tani. 2012. "Imitating Others by Composition of Primitive Actions: A Neuro-Dynamic Model." *Robotics and Autonomous Systems* 60 (5): 729–741. <https://doi.org/10.1016/j.robot.2011.11.005>.
- Bach, Sebastian, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus Robert Müller, and Wojciech Samek. 2015. "On Pixel-Wise Explanations for Non-linear Classifier Decisions by Layer-Wise Relevance Propagation." Edited by Oscar Deniz Suarez. *PLoS One* 10 (7): e0130140. <https://doi.org/10.1371/journal.pone.0130140>.
- Baishya, Shiv S., and Berthold Bäuml. 2016. "Robust Material Classification with a Tactile Skin Using Deep Learning." In *IEEE International Conference on Intelligent Robots and Systems*, 8–15. New York: IEEE. <https://doi.org/10.1109/iros.2016.7758088>.
- Billard, Aude, Sylvain Calinon, Rüdiger Dillmann, and Stefan Schaal. 2008. "Robot Programming by Demonstration." In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1371–1394. Berlin: Springer. https://doi.org/10.1007/978-3-540-30301-5_60.
- Bimbo, Joao, Shan Luo, Kaspar Althoefer, and Hongbin Liu. 2016. "In-Hand Object Pose Estimation Using Covariance-Based Tactile to Geometry Matching." *IEEE Robotics and Automation Letters* 1 (1): 570–577. <https://doi.org/10.1109/lra.2016.2517244>.
- Bousmalis, Konstantinos, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, et al. 2018. "Using Simulation and Domain Adaptation to Improve Efficiency of Deep Robotic Grasping." In *Proceedings—IEEE International Conference on Robotics and Automation*, 4243–4250. New York: IEEE. <https://doi.org/10.1109/icra.2018.8460875>.
- Buckman, Jacob, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. 2018. "Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion." In *Advances in Neural Information Processing Systems*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, 8224–8234. Red Hook, NY: Curran. <http://papers.nips.cc/paper/8044-sample-efficient-reinforcement-learning-with-stochastic-ensemble-value-expansion.pdf>.
- Calandra, Roberto, Andrew Owens, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H. Adelson, and Sergey Levine. 2018. "More than a Feeling: Learning to Grasp and Regrasp Using Vision and Touch." *IEEE Robotics and Automation Letters* 3 (4): 3300–3307. <https://doi.org/10.1109/lra.2018.2852779>.
- Calinon, Sylvain, Florent Guenter, and Aude Billard. 2007. "On Learning, Representing, and Generalizing a Task in a Humanoid Robot." *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 37 (2): 286–298. <https://doi.org/10.1109/tsmcb.2006.886952>.
- Chen, Xi, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. "InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets." In *Advances in Neural Information Processing Systems*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, 2180–2188. Red Hook, NY: Curran. <http://papers.nips.cc/paper/6399-infogan-interpretible-representation-learning-by-information-maximizing-generative-adversarial-nets.pdf>.
- Chen, Yiwen, Shingo Murata, Hiroaki Arie, Tetsuya Ogata, Jun Tani, and Shigeki Sugano. 2016. "Emergence of Interactive Behaviors between Two Robots by Prediction Error Minimization Mechanism." In *2016 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics*, 302–307. New York: IEEE. <https://doi.org/10.1109/devlrm.2016.7846838>.
- Choi, Changhyun, and Henrik I. Christensen. 2012. "3D Pose Estimation of Daily Objects Using an RGB-D Camera." In *IEEE International Conference on Intelligent Robots and Systems*, 3342–3349. New York: IEEE. <https://doi.org/10.1109/iros.2012.6386067>.

- Choi, Changhyun, Yuichi Taguchi, Oncel Tuzel, Ming Yu Liu, and Srikumar Ramalingam. 2012. "Voting-Based Pose Estimation for Robotic Assembly Using a 3D Sensor." In *Proceedings—IEEE International Conference on Robotics and Automation*, 1724–1731. New York: IEEE. <https://doi.org/10.1109/icra.2012.6225371>.
- Dahiya, Ravinder S., Philipp Mittendorf, Maurizio Valle, Gordon Cheng, and Vladimir J. Lumelsky. 2013. "Directions toward Effective Utilization of Tactile Skin: A Review." *IEEE Sensors Journal* 13 (11): 4121–4138. <https://doi.org/10.1109/jssen.2013.2279056>.
- Dong, Siyuan, Wenzhen Yuan, and Edward H. Adelson. 2017. "Improved GelSight Tactile Sensor for Measuring Geometry and Slip." In *IEEE International Conference on Intelligent Robots and Systems*, 137–144. New York: IEEE. <https://doi.org/10.1109/iros.2017.8202149>.
- Duan, Yan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. 2017. "One-Shot Imitation Learning." In *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 1087–1098. Red Hook, NY: Curran. <http://papers.nips.cc/paper/6709-one-shot-imitation-learning.pdf>.
- Ebert, Frederik, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex Lee, and Sergey Levine. 2018. "Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control." ArXiv preprint: <http://arxiv.org/abs/1812.00568>.
- Falco, Pietro, Abdallah Attawia, Matteo Saveriano, and Dongheui Lee. 2018. "On Policy Learning Robust to Irreversible Events: An Application to Robotic In-Hand Manipulation." *IEEE Robotics and Automation Letters* 3 (3): 1482–1489. <https://doi.org/10.1109/lra.2018.2800110>.
- Fang, Kuan, Yunfei Bai, Stefan Hinterstoisser, Silvio Savarese, and Mrinal Kalakrishnan. 2018. "Multi-task Domain Adaptation for Deep Learning of Instance Grasping from Simulation." In *Proceedings—IEEE International Conference on Robotics and Automation*, 3516–3523. New York: IEEE. <https://doi.org/10.1109/icra.2018.8461041>.
- Finn, Chelsea, Pieter Abbeel, and Sergey Levine. 2017. "Model-Agnostic Meta-learning for Fast Adaptation of Deep Networks." In *34th International Conference on Machine Learning, ICML 2017* 3:1856–1868. JMLR.org.
- Finn, Chelsea, Paul Christiano, Pieter Abbeel, and Sergey Levine. 2016. "A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models." ArXiv preprint: <http://arxiv.org/abs/1611.03852>.
- Finn, Chelsea, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. 2017. "One-Shot Visual Imitation Learning via Meta-learning." *Proceedings of the 1st Conference on Robot Learning (CoRL 2017)*, 1–12. ArXiv preprint: <http://arxiv.org/abs/1709.04905>.
- Fishel, Jeremy A., and Gerald E. Loeb. 2012. "Sensing Tactile Microvibrations with the BioTac Comparison with Human Sensitivity." In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechanics*, 1122–1127. New York: IEEE. <https://doi.org/10.1109/biorob.2012.6290741>.
- Friston, Karl J., Jean Daunizeau, James Kilner, and Stefan J. Kiebel. 2010. "Action and Behavior: A Free-Energy Formulation." *Biological Cybernetics* 102 (3): 227–260. <https://doi.org/10.1007/s00422-010-0364-z>.
- Funabashi, Satoshi, Alexander Schmitz, Takashi Sato, Sophon Somlor, and Shigeki Sugano. 2018. "Versatile In-Hand Manipulation of Objects with Different Sizes and Shapes Using Neural Networks." In *IEEE-RAS International Conference on Humanoid Robots*, 768–775. New York: IEEE. <https://doi.org/10.1109/humanoids.2018.8624961>.
- Ganin, Yaroslav, Tejas Kulkarni, Igor Babuschkin, S. M. Ali Eslami, and Oriol Vinyals. 2018. "Synthesizing Programs for Images Using Reinforced Adversarial Learning." ArXiv preprint: <http://arxiv.org/abs/1804.01118>.
- Gao, Yang, Lisa Anne Hendricks, Katherine J. Kuchenbecker, and Trevor Darrell. 2016. "Deep Learning for Tactile Understanding from Visual and Haptic Data." In *Proceedings—IEEE International Conference on Robotics and Automation*, 536–543. New York: IEEE. <https://doi.org/10.1109/icra.2016.7487176>.
- Ha, David, and Jürgen Schmidhuber. 2018. "Recurrent World Models Facilitate Policy Evolution." *Advances in Neural Information Processing Systems* C:2450–2462. ArXiv preprint: <http://arxiv.org/abs/1809.01999>.
- Hafner, Danijar, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. 2018. "Learning Latent Dynamics for Planning from Pixels." ArXiv preprint: <http://arxiv.org/abs/1811.04551>.
- Han, L., Y. S. Guan, Z. X. Li, Q. Shi, and J. C. Trinkle. 1997. "Dextrous Manipulation with Rolling Contacts." In *Proceedings of International Conference on Robotics and Automation 2:992–997*. New York: IEEE. <https://doi.org/10.1109/robot.1997.614264>.
- Han, L., and J. C. Trinkle. 1998. "Dextrous Manipulation by Rolling and Finger Gaiting." In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation* 1:730–735. Cat. No.98CH36146. New York: IEEE. <https://doi.org/10.1109/robot.1998.677060>.
- Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D: Nonlinear Phenomena* 42 (1–3): 335–346. [https://doi.org/10.1016/0167-2789\(90\)90087-6](https://doi.org/10.1016/0167-2789(90)90087-6).

- Heinrich, Stefan, and Stefan Wermter. 2014. “Interactive Language Understanding with Multiple Timescale Recurrent Neural Networks.” In *Lecture Notes in Computer Science*, edited by Stefan Wermter, Cornelius Weber, Włodzisław Duch, Timo Honkela, Petia Koprinkova-Hristova, Sven Magg, Günther Palm, and Alessandro E. P. Villa, 8681 LNCS:193–200. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-11179-7_25.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. “Long Short-Term Memory.” *Neural Computation* 9 (8): 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hodaň, Tomáš, Frank Michel, Eric Brachmann, Wadim Kehl, Anders Glent Buch, Dirk Kraft, Bertram Drost, et al. 2018. “BOP: Benchmark for 6D Object Pose Estimation.” In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11214 LNCS:19–35. https://doi.org/10.1007/978-3-030-01249-6_2.
- Hogan, Francois R., Maria Bauza, Oleguer Canal, Elliott Donlon, and Alberto Rodriguez. 2018. “Tactile Regrasp: Grasp Adjustments via Simulated Tactile Transformations.” In *IEEE International Conference on Intelligent Robots and Systems*, 2963–2970. New York: IEEE. <https://doi.org/10.1109/iros.2018.8593528>.
- Hu, Yinlin, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. 2019. “Segmentation-Driven 6D Object Pose Estimation.” In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 3380–3389. New York: IEEE. <https://doi.org/10.1109/cvpr.2019.00350>.
- Huang, Zhewei, Wen Heng, and Shuchang Zhou. 2019. “Learning to Paint with Model-Based Deep Reinforcement Learning.” ArXiv preprint: <http://arxiv.org/abs/1903.04411>.
- Ijspeert, Auke Jan, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. 2013. “Dynamical Movement Primitives: Learning Attractor Models Formotor Behaviors.” *Neural Computation* 25 (2): 328–373. https://doi.org/10.1162/neco_a_00393.
- Ijspeert, Auke Jan, Jun Nakanishi, and Stefan Schaal. 2002. “Movement Imitation with Nonlinear Dynamical Systems in Humanoid Robots.” In *Proceedings—IEEE International Conference on Robotics and Automation* 2:1398–1403. New York: IEEE. <https://doi.org/10.1109/robot.2002.1014739>.
- Inamura, Tetsunari, Iwaki Toshima, Hiroaki Tanie, and Yoshihiko Nakamura. 2004. “Embodied Symbol Emergence Based on Mimesis Theory.” *International Journal of Robotics Research* 23 (4–5): 363–377. <https://doi.org/10.1177/0278364904042199>.
- Ito, Masato, Kuniaki Noda, Yukiko Hoshino, and Jun Tani. 2006. “Dynamic and Interactive Generation of Object Handling Behaviors by a Small Humanoid Robot Using a Dynamic Neural Network Model.” *Neural Networks* 19 (3): 323–337. <https://doi.org/10.1016/j.neunet.2006.02.007>.
- Iwata, Hiroyasu, and Shigeki Sugano. 2009. “Design of Human Symbiotic Robot TWENDY-ONE.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 580–586. New York: IEEE. <https://doi.org/10.1109/robot.2009.5152702>.
- Johnson, Micah K., and Edward H. Adelson. 2009. “Retrographic Sensing for the Measurement of Surface Texture and Shape.” In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 1070–1077. New York: IEEE. <https://doi.org/10.1109/cvpr.2009.5206534>.
- Jordan, Michael I. 1997. “Serial Order: A Parallel Distributed Processing Approach.” In *Advances in Psychology*, edited by John W. Donahoe and Vivian Packard Dorsel, 121:471–495. Advances in Psychology. Amsterdam: North-Holland. [https://doi.org/10.1016/s0166-4115\(97\)80111-2](https://doi.org/10.1016/s0166-4115(97)80111-2).
- Karl, Maximilian, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt. 2019. “Deep Variational Bayes Filters: Unsupervised Learning of State Space Models from Raw Data.” *5th International Conference on Learning Representations, ICLR 2017—Conference Track Proceedings*. ArXiv preprint: <https://arxiv.org/abs/1605.06432>.
- Kase, Kei, Ryoichi Nakajo, Hiroki Mori, and Tetsuya Ogata. 2019. “Learning Multiple Sensorimotor Units to Complete Compound Tasks Using an RNN with Multiple Attractors.” In *Proceedings of the 2019 IEEE/RSSJ International Conference on Intelligent Robots and Systems*, 4244–4249. <https://doi.org/10.1109/iros40897.2019.8967780>.
- Kase, Kei, Kanata Suzuki, Pin Chu Yang, Hiroki Mori, and Tetsuya Ogata. 2018. “Put-in-Box Task Generated from Multiple Discrete Tasks by a Humanoid Robot Using Deep Learning.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 6447–6452. New York: IEEE. <https://doi.org/10.1109/icra.2018.8460623>.
- Kurutach, Thanard, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. 2018. “Model-Ensemble Trust-Region Policy Optimization.” In *6th International Conference on Learning Representations, ICLR 2018—Conference Track Proceedings*. Available at <https://iclr.cc/Conferences/2018/Schedule>.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. “Deep Learning.” *Nature* 521 (7553): 436–444.
- Lenz, Ian, Honglak Lee, and Ashutosh Saxena. 2015. “Deep Learning for Detecting Robotic Grasps.” *International Journal of Robotics Research* 34 (4–5): 705–724. <https://doi.org/10.1177/0278364914549607>.

- Levine, Sergey, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. 2018. "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection." *International Journal of Robotics Research* 37 (4–5): 421–436. <https://doi.org/10.1177/0278364917710318>.
- Li, Miao, Hang Yin, Kenji Tahara, and Aude Billard. 2014. "Learning Object-Level Impedance Control for Robust Grasping and Dexterous Manipulation." In *Proceedings—IEEE International Conference on Robotics and Automation*, 6784–6791. New York: IEEE. <https://doi.org/10.1109/icra.2014.6907861>.
- Lynch, Corey, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. 2019. "Learning Latent Plans from Play." *Proceedings of the 3rd Conference on Robot Learning (CoRL 2019)*. ArXiv preprint: <http://arxiv.org/abs/1903.01973>.
- Meltzoff, Andrew N., and M. Keith Moore. 1977. "Imitation of Facial and Manual Gestures by Human Neonates." *Science* 198 (4312): 75–78. <https://doi.org/10.1126/science.198.4312.75>.
- Mittendorf, Philipp, and Gordon Cheng. 2011. "Humanoid Multimodal Tactile-Sensing Modules." *IEEE Transactions on Robotics* 27 (3): 401–410. <https://doi.org/10.1109/tro.2011.2106330>.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, et al. 2015. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518 (7540): 529–533. <https://doi.org/10.1038/nature14236>.
- Nagai, Yuki. 2019. "Predictive Learning: Its Key Role in Early Cognitive Development." *Philosophical Transactions of the Royal Society B: Biological Sciences* 374 (1771): 20180030. <https://doi.org/10.1098/rstb.2018.0030>.
- Nair, Ashvin, Dian Chen, Pulkit Agrawal, Phillip Isola, Pieter Abbeel, Jitendra Malik, and Sergey Levine. 2017. "Combining Self-Supervised Learning and Imitation for Vision-Based Rope Manipulation." In *Proceedings—IEEE International Conference on Robotics and Automation*, 2146–2153. New York: IEEE. <https://doi.org/10.1109/icra.2017.7989247>.
- Nakajo, Ryoichi, Shingo Murata, Hiroaki Arie, and Tetsuya Ogata. 2015. "Acquisition of Viewpoint Representation in Imitative Learning from Own Sensory-Motor Experiences." In *5th Joint International Conference on Development and Learning and Epigenetic Robotics, ICDL-EpiRob 2015*, 326–331. New York: IEEE. <https://doi.org/10.1109/devlrm.2015.7346166>.
- Nakajo, Ryoichi, Shingo Murata, Hiroaki Arie, and Tetsuya Ogata. 2018. "Acquisition of Viewpoint Transformation and Action Mappings via Sequence to Sequence Imitative Learning by Deep Neural Networks." *Frontiers in Neurobotics* 12:46. <https://doi.org/10.3389/fnbot.2018.00046>.
- Namikawa, Jun, Ryunosuke Nishimoto, and Jun Tani. 2011. "A Neurodynamic Account of Spontaneous Behavior." *PLoS Computational Biology* 7 (10): e1002221–e1002221. <https://doi.org/10.1371/journal.pcbi.1002221>.
- Ng, Andrew, and Stuart Russell. 2000. "Algorithms for Inverse Reinforcement Learning." In *Proceedings of the Seventeenth International Conference on Machine Learning* 0:663–670. San Francisco: Morgan Kaufmann. <https://doi.org/10.2460/ajvr.67.2.323>.
- Nishihara, Joe, Tomoaki Nakamura, and Takayuki Nagai. 2017. "Online Algorithm for Robots to Learn Object Concepts and Language Model." *IEEE Transactions on Cognitive and Developmental Systems* 9 (3): 255–268. <https://doi.org/10.1109/tcds.2016.2552579>.
- Nishimoto, Ryunosuke, and Jun Tani. 2009. "Development of Hierarchical Structures for Actions and Motor Imagery: A Constructivist View from Synthetic Neuro-Robotics Study." *Psychological Research* 73 (4): 545–558. <https://doi.org/10.1007/s00426-009-0236-0>.
- Noda, Kuniaki, Hiroaki Arie, Yuki Suga, and Tetsuya Ogata. 2014. "Multimodal Integration Learning of Robot Behavior Using Deep Neural Networks." *Robotics and Autonomous Systems* 62 (6): 721–736. <https://doi.org/10.1016/j.robot.2014.03.003>.
- Ogata, Tetsuya, Masamitsu Murase, Jim Tani, Kazunori Komatani, and Hiroshi G. Okuno. 2007. "Two-Way Translation of Compound Sentences and Arm Motions by Recurrent Neural Networks." In *IEEE International Conference on Intelligent Robots and Systems*, 1858–1863. New York: IEEE. <https://doi.org/10.1109/iroso.2007.4399265>.
- Ohmura, Yoshiyuki, Yasuo Kuniyoshi, and Akihiko Nagakubo. 2006. "Conformable and Scalable Tactile Sensor Skin for a Curved Surfaces." In *Proceedings—IEEE International Conference on Robotics and Automation*, 1348–1353. New York: IEEE. <https://doi.org/10.1109/robot.2006.1641896>.
- Pathak, Deepak, Parsa Mahmoudieh, Guanghao Luo, Pulkit Agrawal, Dian Chen, Yide Shentu, Evan Shelhamer, Jitendra Malik, Alexei A. Efros, and Trevor Darrell. 2018. "Zero-Shot Visual Imitation." In *6th International Conference on Learning Representations, ICLR 2018—Conference Track Proceedings*, 2050–2053. Available at <https://iclr.cc/Conferences/2018/Schedule>.
- Paulino, Tiago, Pedro Ribeiro, Miguel Neto, Susana Cardoso, Alexander Schmitz, Jose Santos-Victor, Alexandre Bernardino, and Lorenzo Jamone. 2017. "Low-Cost 3-Axis Soft Tactile Sensors for the Human-Friendly Robot

- Vizy.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 966–971. New York: IEEE. <https://doi.org/10.1109/icra.2017.7989118>.
- Peng, Xue Bin, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. “DeepMimic: Example-Guided Deep Reinforcement Learning of Physics-Based Character Skills.” *ACM Transactions on Graphics* 37 (4): 1–14. <https://doi.org/10.1145/3197517.3201311>.
- Peng, Xue Bin, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. 2018. “SFV: Reinforcement Learning of Physical Skills from Videos” 37 (6). <https://doi.org/10.1145/3272127.3275014>.
- Peng, Xue Bin, Angjoo Kanazawa, Sam Toyer, Pieter Abbeel, and Sergey Levine. 2018. “Variational Discriminator Bottleneck: Improving Imitation Learning, Inverse RL, and GANs by Constraining Information Flow.” ArXiv preprint: <http://arxiv.org/abs/1810.00821>.
- Redmon, Joseph, and Anelia Angelova. 2015. “Real-Time Grasp Detection Using Convolutional Neural Networks.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 1316–1322. New York: IEEE. <https://doi.org/10.1109/icra.2015.7139361>.
- Rusinkiewicz, Szymon, and Marc Levoy. 2001. “Efficient Variants of the ICP Algorithm.” In *Proceedings of International Conference on 3-D Digital Imaging and Modeling, 3DIM*, 145–152. New York: IEEE. <https://doi.org/10.1109/im.2001.924423>.
- Schmitz, Alexander, Yusuke Bansho, Kuniaki Noda, Hiroyasu Iwata, Tetsuya Ogata, and Shigeki Sugano. 2014. “Tactile Object Recognition Using Deep Learning and Dropout.” In *IEEE-RAS International Conference on Humanoid Robots*, 1044–1050. New York: IEEE. <https://doi.org/10.1109/humanoids.2014.7041493>.
- Schwarz, Max, Hannes Schulz, and Sven Behnke. 2015. “RGB-D Object Recognition and Pose Estimation Based on Pre-trained Convolutional Neural Network Features.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 1329–1335. New York: IEEE. <https://doi.org/10.1109/icra.2015.7139363>.
- Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, et al. 2018. “A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play.” *Science* 362 (6419): 1140–1144. <https://doi.org/10.1126/science.aar6404>.
- Smilkov, Daniel, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. 2017. “SmoothGrad: Removing Noise by Adding Noise.” ArXiv preprint: 1706.03825. <http://arxiv.org/abs/1706.03825>.
- Stramandinoli, Francesca, Davide Marocco, and Angelo Cangelosi. 2017. “Making Sense of Words: A Robotic Model for Language Abstraction.” *Autonomous Robots* 41 (2): 367–383. <https://doi.org/10.1007/s10514-016-9587-8>.
- Sugita, Yuuya, and Jun Tani. 2005. “Learning Semantic Combinatoriality from the Interaction between Linguistic and Behavioral Processes.” *Adaptive Behavior* 13 (1): 33–52. <https://doi.org/10.1177/105971230501300102>.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. 2014. “Sequence to Sequence Learning with Neural Networks.” In *Advances in Neural Information Processing Systems* 4:3104–3112.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Takahashi, Kuniyuki, and Jethro Tan. 2019. “Deep Visuo-tactile Learning: Estimation of Tactile Properties from Images.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 8951–8957. New York: IEEE. <https://doi.org/10.1109/icra.2019.8794285>.
- Tellex, Stefanie, Thomas Kollar, Steven Dickerson, Matthew R. Walter, Ashis Gopal Banerjee, Seth Teller, and Nicholas Roy. 2011. “Understanding Natural Language Commands for Robotic Navigation and Mobile Manipulation.” In *Proceedings of the National Conference on Artificial Intelligence* 2:1507–1514.
- Tian, Stephen, Frederik Ebert, Dinesh Jayaraman, Mayur Mudigonda, Chelsea Finn, Roberto Calandra, and Sergey Levine. 2019. “Manipulation by Feel: Touch-Based Control with Deep Predictive Models.” In *Proceedings—IEEE International Conference on Robotics and Automation*, 818–824. New York: IEEE. <https://doi.org/10.1109/icra.2019.8794219>.
- Tobin, Josh, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. “Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World.” In *IEEE International Conference on Intelligent Robots and Systems*. New York: IEEE. <https://doi.org/10.1109/iros.2017.8202133>.
- Tomo, Tito Pradhono, Alexander Schmitz, Wai Keat Wong, Harris Kristanto, Sophon Somlor, Jinsun Hwang, Lorenzo Jamone, and Shigeki Sugano. 2018. “Covering a Robot Fingertip with USkin: A Soft Electronic Skin with Distributed 3-Axis Force Sensitive Elements for Robot Hands.” *IEEE Robotics and Automation Letters* 3 (1): 124–131. <https://doi.org/10.1109/lra.2017.2734965>.
- Vinyals, Oriol, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, et al. 2019. “Grandmaster Level in StarCraft II Using Multi-agent Reinforcement Learning.” *Nature* 575 (7782): 350–354. <https://doi.org/10.1038/s41586-019-1724-z>.
- Wu, Bohan, Ireteayo Akinola, Jacob Varley, and Peter Allen. 2019. “MAT: Multi-fingered Adaptive Tactile Grasping via Deep Reinforcement Learning.” *3rd Conference on Robot Learning*. ArXiv preprint: <https://arxiv.org/abs/1909.04787>.

- Xu, Kelvin, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention." In *32nd International Conference on Machine Learning 2015* 3:2048–2057.
- Yamada, Tatsuro, Hiroyuki Matsunaga, and Tetsuya Ogata. 2018. "Paired Recurrent Autoencoders for Bidirectional Translation between Robot Actions and Linguistic Descriptions." *IEEE Robotics and Automation Letters* 3 (4): 3441–3448. <https://doi.org/10.1109/lra.2018.2852838>.
- Yamaguchi, Akihiko, and Christopher G. Atkeson. 2016. "Combining Finger Vision and Optical Tactile Sensing: Reducing and Handling Errors While Cutting Vegetables." In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 1045–1051. New York: IEEE. <https://doi.org/10.1109/humanoids.2016.7803400>.
- Yamashita, Yuichi, and Jun Tani. 2008. "Emergence of Functional Hierarchy in a Multiple Timescale Neural Network Model: A Humanoid Robot Experiment." *PLoS Computational Biology* 4 (11). <https://doi.org/10.1371/journal.pcbi.1000220>.
- Yang, Haolin, Fuchun Sun, Wenbing Huang, Lele Cao, and Bin Fang. 2016. "Tactile Sequence Based Object Categorization: A Bag of Features Modeled by Linear Dynamic System with Symmetric Transition Matrix." In *Proceedings of the International Joint Conference on Neural Networks*, 5218–5225. New York: IEEE. <https://doi.org/10.1109/ijcnn.2016.7727889>.
- Yang, Pin Chu, Kazuma Sasaki, Kanata Suzuki, Kei Kase, Shigeki Sugano, and Tetsuya Ogata. 2017. "Repeatable Folding Task by Humanoid Robot Worker Using Deep Learning." *IEEE Robotics and Automation Letters* 2 (2): 397–403. <https://doi.org/10.1109/lra.2016.2633383>.
- Yang, Yezhou, Yi Li, Cornelia Fermüller, and Yiannis Aloimonos. 2015. "Robot Learning Manipulation Action Plans by 'Watching' Unconstrained Videos from the World Wide Web." In *Proceedings of the National Conference on Artificial Intelligence* 5:3686–3692.
- Yu, Tianhe, Chelsea Finn, Annie Xie, Sudeep Dasari, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. 2018. "One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-learning." In *Proceedings of the Robotics: Science and Systems XIV (RSS 2018)*, 1–12. <https://doi.org/10.15607/rss.2018.xiv.002>.
- Yuan, Wenzhen, Siyuan Dong, and Edward H. Adelson. 2017. "GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force." *Sensors* 17 (12): 2762. <https://doi.org/10.3390/s17122762>.
- Yuan, Wenzhen, Shaoxiong Wang, Siyuan Dong, and Edward Adelson. 2017. "Connecting Look and Feel: Associating the Visual and Tactile Properties of Physical Materials." In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, 4494–4502. New York: IEEE. <https://doi.org/10.1109/cvpr.2017.478>.
- Yuan, Wenzhen, Chenzhuo Zhu, Andrew Owens, Mandayam A. Srinivasan, and Edward H. Adelson. 2017. "Shape-Independent Hardness Estimation Using Deep Learning and a GelSight Tactile Sensor." In *Proceedings—IEEE International Conference on Robotics and Automation*, 951–958. New York: IEEE. <https://doi.org/10.1109/icra.2017.7989116>.
- Zhang, Yazhan, Weihao Yuan, Zicheng Kan, and Michael Yu Wang. 2020. "Towards Learning to Detect and Predict Contact Events on Vision-Based Tactile Sensors." *3rd Conference on Robot Learning*, 1395–1404. ArXiv preprint: <https://arxiv.org/abs/1910.03973>.

10 Cognitive Architectures

David Vernon

10.1 Introduction

As the definition of cognitive robotics in chapter 1 makes clear, the field draws on several disciplines, including robotics, artificial intelligence, and cognitive science. Its goal is to design an integrated cognitive system that combines a range of abilities, such as sensorimotor behaviors, knowledge-based reasoning, and social skills, in the form of an intelligent robot. Its foundations in systems engineering and cognitive science coalesce in a single concept: a cognitive architecture.

From the perspective of systems engineering, a cognitive architecture mirrors the system architecture, using the power of abstraction to render the modeling, specification, and design of a complete complex system tractable.

From the perspective of cognitive science, in which the term “cognitive architecture” originates (Newell 1990), the concept of a cognitive architecture is the result of over sixty years of research. To understand what it means from this perspective requires us to first familiarize ourselves with the roots of cognitive science and the different paradigms that exist within that discipline. In turn, this will allow us to understand the different types of cognitive architecture and the role a cognitive architecture plays in cognitive science, in general, and cognitive robotics, in particular.

With this understanding in place, we review the key attributes of a cognitive architecture before surveying the core cognitive abilities of the many cognitive architectures that exist today. We examine two cognitive architectures in some detail to highlight the way these abilities are realized in cognitive robots. We finish by exploring what the future might hold for cognitive architectures and the challenges that remain.

10.2 The Foundations of Cognitive Science

Cognitive science embraces neuroscience, cognitive psychology, linguistics, epistemology, philosophy, and artificial intelligence, among other disciplines. Its primary goal is to explain the underlying processes of human cognition, ideally in the form of a model that can be replicated in artificial agents. It has its roots in cybernetics in the early 1940s (Wiener 1948) but appears as a formal discipline referred to as cognitivism in the late

1950s. Cognitivism built on the logical foundations laid by the early cyberneticians and exploited the computer as a literal metaphor for cognitive function and operation, using symbolic information processing as its core model of cognition. Cybernetics also gave rise to the alternative emergent systems approach, which recognized the importance of self-organization in cognitive processes, eventually embracing connectionism, dynamical systems theory, and the enactive perspective on cognitive science. Hybrid systems attempt to combine the cognitivist and emergent paradigms to varying degrees, quite often ignoring some of the incompatible assumptions that the cognitivist and the emergent paradigms make about the fundamental nature of cognition (Vernon 2014).

10.2.1 The Cognitivist Paradigm of Cognitive Science

The cognitivist paradigm, which embraces artificial intelligence (AI), dates from a conference held at Dartmouth College, New Hampshire, in July and August 1956. It was attended by Allen Newell, Herbert Simon, John McCarthy, Marvin Minsky, and Claude Shannon, among others, all of whom exerted a very significant influence on the development of AI over the next half century.

The essential position of cognitivism is that cognition is achieved by computations performed on internal symbolic knowledge representations in a process whereby information about the world is taken in through the senses, filtered by perceptual processes to generate descriptions that abstract away irrelevant data, represented in symbolic form, and reasoned about to infer what is required to perform some task and achieve some goal. In the cognitivist paradigm, any physical platform that supports the performance of the required symbolic computation will suffice. In other words, the physical realization of the computational model is inconsequential to the model. The principled decoupling of computational operation from the physical platform that supports these computations is referred to as computational functionalism (Piccinini 2010). Allen Newell made several landmark contributions to the establishment of practical cognitivist systems: in the early 1980s with his introduction of the concept of a knowledge-level system, the maximum rationality hypothesis, and the principle of rationality (Newell 1982); in the mid-1980s with the development of the Soar cognitive architecture for general intelligence (along with John Laird and Paul Rosenbloom; Laird, Newell, and Rosenbloom 1987); and in 1990 with the concept of a unified theory of cognition (Newell 1990)—that is, a theory that covers a broad range of cognitive issues, such as attention, memory, problem-solving, decision-making, and learning from several aspects, including psychology, neuroscience, and computer science.

10.2.2 The Emergent Paradigm of Cognitive Science

In the emergent paradigm, cognition is one of the processes by which an autonomous system maintains its autonomy. Through cognition, the system constructs its reality—its world and the meaning of its perceptions and actions—as a result of its operation in that world. This process of making sense of its environmental interactions is one of the foundations of a branch of cognitive science called *enaction* (Stewart, Gapenne, and Di Paolo 2010; Vernon 2010). Cognition is also the means by which the system prepares for interaction that may be necessary in the future. Thus, cognition is intrinsically linked with the ability of an agent to act prospectively. As such, many emergent approaches focus on the acquisition of anticipatory skills rather than knowledge, asserting that processes that guide

action and improve the capacity to guide action form the root capacity of all intelligent systems (Christensen and Hooker 2000). As a result, in contrast to cognitivism, emergent approaches are necessarily embodied, and the physical form of the agent's body plays a causal role in the cognitive process. Together, the body and the brain form the basis of a cognitive system, and they do so in the context of a structured environmental niche to which the body is adapted. Because of this, cognition in the emergent paradigm is sometimes referred to as *embodied cognition*, although some emergent approaches make even stronger assertions about the nature of cognition. The emergent paradigm typically exploits connectionism or dynamical systems theory. In general, connectionist systems correspond to models at a lower level of abstraction, dynamical systems to a higher level. They are sometimes referred to as subsymbolic processes.

10.2.3 Hybrid Systems

Hybrid systems are attempts to exploit both the cognitivist and emergent paradigms of cognitive science. They exploit symbolic knowledge to represent the agent's world and logical rule-based systems to reason with this knowledge to pursue tasks and achieve goals. At the same time, they typically use emergent models of perception and action to explore the world and construct this knowledge. Hybrid systems use both symbolic and subsymbolic representations. The latter are constructed using subsymbolic connectionist processes as the system interacts with and explores the world. So, instead of a designer programming in all the necessary knowledge, objects and events in the world can be represented by observed correspondences between sensed perceptions, agent actions, and sensed outcomes. Thus, as with an emergent system, a hybrid system's ability to understand the external world is dependent on its ability to flexibly interact with it. Interaction becomes an organizing mechanism that establishes a learned association between perception and action. For a detailed comparison of cognitivist, emergent, and hybrid paradigms of cognitive science, see Vernon, Metta, and Sandini (2007b) and Vernon (2014).

10.3 The Types of Cognitive Architecture

A cognitive architecture is a software framework that integrates all the elements required for a system to exhibit the attributes considered to be characteristic of a cognitive agent. Just what these elements are is open to interpretation, but as we will see, there is common ground in the identification of core cognitive abilities in these interpretations—for example, perception, action, learning, adaptation, anticipation, motivation, autonomy, internal simulation, attention, action selection, memory, reasoning, and metareasoning (Vernon 2014; Vernon, von Hofsten, and Fadiga 2016; Kotseruba and Tsotsos 2020).

Furthermore, a cognitive architecture determines the overall structure and organization of a cognitive system, including the component parts or modules (Sun 2004), the relations between these modules, and the essential algorithmic and representational details within them (Langley 2006). The architecture specifies the formalisms for knowledge representations and the types of memories used to store them, the processes that act upon that knowledge, and the learning mechanisms that acquire it. For cognitivist and hybrid approaches, a cognitive architecture also provides a way of programming the system so that domain and task knowledge can be embedded in the system.

A cognitive architecture makes explicit the set of assumptions upon which that cognitive model is founded. These assumptions are typically derived from several sources: biological or psychological data, philosophical arguments, or working hypotheses inspired by work in different disciplines such as neurophysiology, psychology, and artificial intelligence. Once it has been created, a cognitive architecture also provides a framework for developing the ideas and assumptions encapsulated in the architecture.

There are three different types of cognitive architecture, each derived from the three paradigms of cognitive science: the cognitivist, the emergent, and the hybrid. Cognitivist cognitive architectures are often referred to as symbolic cognitive architectures (Kotseruba and Tsotsos 2020). It is noteworthy that the term “cognitive architecture” itself is due to Allen Newell and his colleagues in their work on unified theories of cognition (Newell 1990). Consequently, for cognitivism a cognitive architecture represents any attempt to create a unified theory of cognition. The cognitive architectures Soar (Laird, Newell, and Rosenbloom 1987; Laird 2009, 2012), ACT-R (Anderson 1996; Anderson et al. 2004), and CLARION (Sun 2007, 2016) are archetypal candidate unified theories of cognition, all of which are classified as hybrid cognitive architectures in the survey by Kotseruba and Tsotsos (2020).

10.3.1 The Cognitivist Perspective on Cognitive Architecture

In the cognitivist paradigm, the focus in a cognitive architecture is on the aspects of cognition that are constant over time and that are independent of the task (Ritter and Young 2001; Langley, Laird, and Rogers 2009). A cognitivist cognitive architecture is a generic computational model that is neither domain-specific nor task-specific, and it needs to be provided with knowledge to perform any given task. This combination of a given cognitive architecture and a particular knowledge set is generally referred to as a *cognitive model*. In many cognitivist systems, much of the knowledge incorporated into the model is normally provided by the designer, and often this knowledge is highly crafted, possibly drawing on years of experience working in the problem domain. Machine learning is increasingly used to augment and adapt this knowledge.

10.3.2 The Emergent Perspective on Cognitive Architecture

Emergent approaches to cognition focus on the development of the agent from a primitive state to a fully cognitive state over its lifetime. As such, an emergent cognitive architecture is the initial state from which an agent subsequently develops. Development requires exposure to an environment that is conducive to development, one in which there is sufficient regularity to allow the system to build a sense of understanding of the world around it but not excessive variety that would overwhelm an agent that has inherent limitations on the speed with which it can develop. Thus, emergent cognition has two aspects, architecture and gradually acquired experience, mirroring the two aspects of a cognitivist cognitive architecture: architecture and knowledge. These two aspects of emergent cognition are referred to as phylogeny and ontogeny (or ontogenesis), the latter being the interactions and experiences that a developing cognitive system is exposed to as it acquires an increasing degree of cognitive capability. Since the emergent paradigm holds that the physical system—the body—is also a part of the cognitive process, an emergent cognitive archi-

ecture should reflect in some way the structure and capabilities—the morphology—of the physical body in which it is embedded.

10.3.3 The Hybrid Perspective on Cognitive Architecture

As we have noted, hybrid systems endeavor to have the best of both worlds, combining the strengths of the cognitivist and emergent approaches. Most hybrid systems focus on integrating symbolic and subsymbolic (usually connectionist) processing.

Hybrid cognitive architectures are the most prevalent type. The survey by Kotseruba and Tsotsos (2020) lists twenty-two symbolic (i.e., cognitivist) cognitive architectures, fourteen emergent, and forty-eight hybrid, thirty-eight of which are fully integrated.

10.4 Desirable Characteristics of a Cognitive Architecture

If a cognitive architecture is intended to be a unified theory of cognition, as most cognitivist cognitive architectures are, then it should exhibit certain desirable attributes—*desiderata*—including ecological realism, bioevolutionary realism, cognitive realism, and eclecticism of methodologies and techniques, as well as several behavioral characteristics (Sun 2004). Ecological realism means that a cognitive architecture should focus on allowing the cognitive system to operate in its natural environment, engaging in everyday activities and dealing with many concurrent and often conflicting goals with many environmental contingencies. Bioevolutionary realism means that a cognitive model of human intelligence should be reducible to a model of animal intelligence. Cognitive realism means that a cognitive architecture should attempt to capture the essential characteristics of human cognition from the perspectives of psychology, neuroscience, and philosophy. Eclecticism of methodologies and techniques means that new models should draw on, subsume, or supersede older models. Most cognitive architectures for cognitive robotics are not intended to be a unified theory of cognition, and consequently, these attributes can be addressed only to the extent that they are useful from a robotics perspective.

In the emergent paradigm of cognitive science, development is the process whereby a cognitive agent 1) expands its repertoire of action capabilities and 2) extends the time horizon of its ability to anticipate events in its world, including the need to act, the outcome of selected actions, the intentions of other cognitive agents, and the outcome of their actions (Vernon 2010). These considerations give rise to an additional set of *desiderata* for developmental cognitive architectures (Vernon, von Hofsten, and Fadiga 2016), including the need for a value system to determine the goals of actions and provide the drive for achieving them (Oudeyer, Kaplan, and Hafner 2007; Merrick 2010) along with exploratory and social motives (Piaget 1954; Vygotsky 1978; Lindblom 2015) to modulate behavior and select actions (Edelman 2006). The adaptation inherent in development is dependent on learning. A developmental cognitive architecture needs to have at least three different modes of learning: supervised learning, reinforcement learning, and unsupervised learning (Doya 1999). It also requires some mechanism to simulate future events (Seligman et al. 2013), to simulate the execution of actions and the likely outcome of those actions (Hesslow 2002, 2012), and to take alternative perspectives, including those of other agents (Schacter, Addis, and Buckner 2008).

10.5 Surveys of Cognitive Architectures

While several surveys of cognitive architectures have been published over the past ten or so years (Vernon, Metta, and Sandini 2007b; Duch, Oentaryo, and Pasquier 2008; Samsanovich 2010; Thórisson and Helgasson 2012), the recent survey by Kotseruba and Tsotsos (2020) is by far the most comprehensive. It targets eighty-four cognitive architectures, estimating that approximately three hundred cognitive architectures have been developed and that approximately one-third are currently active. The most often cited cognitive architectures are ACT-R (Anderson 1996; Anderson et al. 2004), Soar (Laird, Newell, and Rosenbloom 1987; Laird 2012), CLARION (Sun 2007, 2016), ICARUS (Langley 2006; Langley and Choi 1999 [2006]), EPIC (Kieras and Meyer 1997), and LIDA (Franklin et al. 2007, 2014). The majority of cognitive architectures focus on modeling human cognition.

Despite its comprehensive coverage, almost inevitably the Kotseruba and Tsotsos survey is not complete. For example, it omits the CRAM cognitive architecture (Beetz et al. 2010; Mösenlechner 2016), possibly because the CRAM literature refers to a cognitive robot abstract machine and to cognition-enabled robotics, rather than a cognitive architecture. Later in the chapter, we use CRAM as one of our two examples of cognitive architectures for cognitive robotics. Nevertheless, the survey provides a peerless basis on which to compare and contrast existing cognitive architectures by addressing the extent to which they exhibit core cognitive abilities, and we will refer to it throughout this section.

10.5.1 Comparing Cognitive Architectures

Despite efforts to establish an agreed set of criteria for comparing and evaluating cognitive architectures based on desirable characteristics such as Sun's (2004) desiderata and Newell's (1990, 1992) functional criteria, disagreements persist regarding the research goals, structure, operation, and application of cognitive architectures. Because of this, and in the absence of a clear definition and general theory of cognition, not to mention difficulties in defining intelligence, Kotseruba and Tsotsos adopt a pragmatic approach, treating intelligence as a set of system competences and behaviors. Thus, rather than summarize and review each cognitive architecture individually, Kotseruba and Tsotsos address seven core cognitive abilities—perception, attention mechanisms, action selection, memory, learning, reasoning, and metareasoning—and discuss the degree to which the eighty-four architectures surveyed exhibit these abilities. Significantly, they don't include anticipation (i.e., prospection) as a core cognitive ability as others do, both in surveys of cognitive architectures (Vernon, Metta, and Sandini 2007b) or in the cognitive science literature (Atance and O'Neill 2001; Gilbert and Wilson 2007; Schacter et al. 2012; Seligman et al. 2013). On the other hand, they do include attention, reasoning, and metacognition, three pivotal abilities that have often been omitted from other surveys. We summarize these seven core cognitive abilities in the following, adding, for completeness, a short note on the central role played by anticipation (i.e., prospection) in cognition and cognitive architectures.

10.5.2 Core Cognitive Abilities

Perception

Perception is a process that transforms raw input into the system's internal representation. Vision is the most commonly implemented sensory modality, but more than half of the

cognitive architectures surveyed use simulated visual input rather than transforming the raw sensory data. In general, symbolic cognitive architectures tend to have limited perceptual abilities, and therefore they rely on direct simulated data input. Audition is less commonly found in cognitive architectures, while touch, smell, and proprioception are rarely implemented with any fidelity. Most architectures use only two modalities simultaneously—for example, vision and audition or vision and range data (e.g., from Lidar sensors). Only a few architectures aim for biological fidelity in perception. For the most part, cognitive architectures ignore crossmodal interaction and adopt a modular approach when dealing with sensory modalities, despite its importance in developmental robotics (Cangelosi and Schlesinger 2015).

Attention

Attention is a process that reduces the information a cognitive system has to process, selecting relevant information and filtering out irrelevant information from sensory data. Kotseruba and Tsotsos refer to three classes of information reduction mechanism (Tsotsos 2011): *selection*, *restriction*, and *suppression*. Selection mechanisms choose one entity from many—for example, gaze and viewpoint selection, restriction mechanisms choose some entities from many, and suppression mechanisms suppress some entities from many. The restrictive mechanism reduces the search space by priming—that is, preparing the visual system for input based on task requirements, exogenous motivations (e.g., domain knowledge), exogenous cues (external stimuli), exogenous tasks (restricting attention to objects relevant to the task), and visual field (limiting the visual field). The suppression mechanisms include feature or spatial inhibition, task-irrelevant stimuli suppression, negative priming, and location or object inhibition of return to bias the agent returning attention to previously attended locations. The most frequently implemented mechanisms of attention are selection and restriction, with only a few cognitive architectures implementing a suppression mechanism. Kotseruba and Tsotsos note that visual attention is largely overlooked in cognitive architectures, with exceptions including the ISAC (Kawamura et al. 2008) and iCub cognitive architectures (Vernon, Sandini, and Metta 2007a).

Action Selection

Action selection determines what the agent should do next. There are two major approaches: planning and dynamic action selection. Planning, using traditional AI techniques, determines a sequence of steps to reach a certain goal or solve a problem prior to execution of the plan. Dynamic action selection involves the selection of one action based on knowledge at the time, typically using winner-take-all, probabilistic, or predefined order selection mechanisms. The criteria for selection include relevance, utility (in the sense of expected contribution to the current goal), and internal functions—for example, transient emotion, drives, or internal mechanisms, including basic physiological needs and high-level social drives and personality traits that bias or modulate the action selection rather than directly determining the next behavior. Planning, prevalent in symbolic architectures and in hybrid architectures but also found in emergent architectures, is often augmented with dynamic action selection mechanisms to improve the capability for adaptivity to environmental changes.

Memory

Kotseruba and Tsotsos identify six types of memory in cognitive architectures: short-term sensory memory and working memory and long-term episodic, semantic, procedural, and

global memory. Sensory memory is a very short-term buffer that stores several recent percepts and has a decay rate in the region of tens of milliseconds for visual data, longer for aural data. Working memory is temporary storage for percepts and information related to the current task and is frequently associated with the current focus of attention. It is critical for attention, reasoning, and learning.

Episodic memory (Tulving 1972, 1984) plays a key role in the anticipatory aspect of cognition. It refers to specific instances in the agent's experience, while semantic memory refers to general knowledge about the agent's world that may be independent of the agent's specific experience: knowledge of general facts about objects and concepts and the relationships between those objects. In symbolic cognitive architectures, semantic memory is often represented as a graph-like ontology network, the nodes being the concepts and the links the relationships. In emergent cognitive architectures, semantic memory is typically represented by a pattern of activity in a connectionist network.

Episodic and semantic memory are collectively known as declarative memory. Declarative memory captures knowledge, while procedural memory captures skills, equipping an agent to "know that" and "know how," respectively (Ryle 1949).

In symbolic production systems, procedural knowledge is the knowledge of how to carry out some task, represented by a set of if-then rules preprogrammed or learned for a particular domain. In emergent systems, procedural memory may comprise sequences of state-action pairs or perceptuomotor associations.

Global memory is reserved for cognitive architectures that don't draw the type-duration distinction and use a unified global structure for all knowledge.

Learning

Learning refers to an ability for a system to improve its performance over time through the acquisition of knowledge or skill. Two types of learning can be distinguished: declarative and nondeclarative. Declarative learning is concerned with explicit knowledge acquisition, while nondeclarative learning focuses on perceptual, procedural, associative, and nonassociative learning.

Of the eighty-four cognitive architectures surveyed by Kotseruba and Tsotsos, nineteen—mostly symbolic and hybrid—do not implement learning of any type.

Declarative learning can take several forms. In production systems, new declarative knowledge—facts about the world—are learned when either a fact or a rule is added to declarative memory—for example, after completing a goal or resolving an impasse. Thus, new symbolic knowledge is learned when local inference rules are applied to existing knowledge to obtain new knowledge, encapsulated in what is referred to as a *chunk*. In emergent and hybrid cognitive architectures, declarative learning often takes the form of the association of perceptual features with the identity of objects.

Perceptual learning refers to learning about the environment from perceptual data: uncovering perceptual patterns, constructing associations between percepts, and inferring knowledge about the environment—for example, its spatial organization.

Procedural learning refers to learning skills by repetitive practice until the skill becomes automatic. Note that this view of procedural learning entails a different view of what constitutes procedural knowledge compared with procedural knowledge in cognitivist production systems.

Associative learning is used to refer to the process of improving decision-making through the influence of reward and punishment. Reinforcement learning is often used as a computational model of associative learning, including variants such as temporal difference learning, Q-learning, and Hebbian learning. Nearly half the cognitive architectures surveyed use reinforcement learning to implement associative learning. Since reinforcement learning can be used with many forms of representation, it is used in all types of cognitive architecture: symbolic, emergent, and hybrid. In symbolic (i.e., cognitivist) cognitive architectures, reinforcement learning facilitates adaptation by weighting the importance of beliefs and actions based on the outcome of their use. In hybrid and emergent cognitive architectures, reinforcement learning also facilitates adaptation, but in these cases by establishing weighted associations between states and actions. This often takes the form of an initial phase of motor babbling—that is, performing random movements and observing their sensory outcome, followed by a learning phase to establish stable patterns known as sensorimotor contingencies.

Nonassociative learning refers to an often gradual adjustment of the weighting or importance of a single system entity, rather than an associative linking between two or more entities—for example, the gradual reduction of the strength of a response to some stimulus or pattern of system activity that is repeatedly presented. This is known as habituation. Sensitization has the opposite effect, such as a gradual increase in the strength of response to some repeated stimulus or activity.

Kotseruba and Tsotsos note that, surprisingly, deep learning does not yet feature strongly in cognitive architectures, but it is likely to play an important role in the future. We return to this topic in section 10.7.

Reasoning

Reasoning is the ability to logically and systematically process knowledge, typically to infer conclusions. The three classical forms of logical inference are deduction, induction, and abduction. In the context of cognitive architectures, reasoning focusses on the practical objective of finding the next (best) action to perform. Cognitive architectures typically aim to facilitate human-level intelligence, but they do not necessarily try to model the processes of human reasoning. Those that do include ACT-R (Anderson 1996; Anderson et al. 2004), Soar (Laird et al. 1987; Laird 2009, 2012), and CLARION (Sun 2007, 2016). Many emergent cognitive architectures do not address reasoning, even if they are capable of facilitating complex behavior. Some emergent cognitive architectures, such as SPA (Eliasmith et al. 2012), effect symbolic reasoning using neural architectures, raising the possibility that it might not be necessary to introduce a hard distinction between symbolic cognition and subsymbolic cognition.

Metacognition

Metacognition refers to a cognitive system's ability to monitor its internal cognitive processes and reason about them, acquiring data about the internal operation and status of the cognitive system—for example, availability of internal resources, confidence values during task execution, and sometimes generating temporal traces of activity during task execution. Approximately one-third of the eighty-four cognitive architectures surveyed by Kotseruba and Tsotsos have a metacognition element. These are mainly symbolic cognitive

architectures and hybrid cognitive architectures with a strong component of symbolic processing. Metacognition is needed for social cognition, especially if the cognitive architecture is to form a *theory of mind*, also known as perspective taking—that is, the ability to infer the cognitive states of other agents with which it is interacting, predicting their behavior, and acting appropriately. Very few cognitive architectures support this ability. Kotseruba and Tsotsos note only two: Sigma (Rosenbloom, Demski, and Ustun 2016) and PolyScheme (Trafton et al. 2005).

Prospection

Although the core cognitive abilities identified by Kotseruba and Tsotsos do not include prospection, it plays such a central role in cognition that we include it here for completeness.

Prospection—the capacity to anticipate the future—is one of the hallmark attributes of cognition. It also lies at the heart of the other core characteristics of a cognitive agent: autonomy, perception, action, learning, and adaptation (Vernon 2014). It facilitates autonomy and the ability to cope with adversarial conditions by allowing the agent to prepare to act. It is also involved in constitutive autonomy (Froese, Virgo, and Izquierdo 2007), predictively adjusting internal system processes through allostasis (Sterling 2012). It facilitates perception through expectation-driven attentional processes (Borji, Sihite, and Itti 2014). Attention, in turn, facilitates predictive control of, for example, gaze (Flanagan and Johansson 2003) and the prediction of the consequences of actions (Flanagan et al. 2013). In general, anticipation is central to action since actions are goal directed and guided by prospective information (von Hofsten 2009): a cognitive agent continually anticipates the need to act, and it anticipates the outcome of those actions (Vernon, von Hofsten, and Fadiga 2011). Prospection also lies at the heart of learning, for learned models are used both for prediction and explanation. Finally, adaptivity arises in cognitive agents when the learned models fail to produce accurate or reliable predictions.

Consensus is emerging that internal simulation plays a key role in prospection (Svensson, Lindblom, and Ziemke 2007; Mohan, Bhat, and Morasso 2018). However, there is less agreement about the manner in which internal simulation is accomplished. Some cognitive architectures opt for an explicit module in the architecture (e.g., Kawamura et al. 2008; Beetz et al. 2010; Kunze and Beetz 2017), while in others it is a covert mode of operation, with internal simulation effected by the same subsystems as those responsible for sensorimotor-mediated action but using covert, internally generated endogenous sensorimotor signals rather than exogenous sensorimotor signals (e.g., Demiris and Khadhouri 2006; Shanahan 2006).

10.5.3 Applications

Kotseruba and Tsotsos identify more than nine hundred projects that use one of the eighty-four cognitive architectures surveyed. They identify ten classes of application: psychological experiments, robotics, human performance modeling, human-robot and human-computer interaction, natural language processing, categorization and clustering, computer vision, games and puzzles, virtual agents, and miscellaneous projects that don't fall into the other nine classes. Robotics applications account for nearly a quarter of all applications, mainly for navigation and obstacle avoidance, fetch and carry tasks, object localization, and object manipulation.

10.6 Example Cognitive Architectures

To highlight the issues we have covered so far, in this section we examine two examples of cognitive architectures that focus specifically on cognitive robotics: CRAM (Betz, Mösenlechner, and Tenorth 2010; Mösenlechner 2016), a knowledge-based reasoning architecture, and ISAC (Kawamura et al. 2008), an architecture built from communicating software agents and memory subsystems.

10.6.1 The CRAM Cognitive Architecture

CRAM stands for cognitive robot abstract machine. It is a hybrid cognitive architecture, first introduced in 2010 (Betz, Mösenlechner, and Tenorth 2010). Since then it has developed significantly, building on the original basis for the architecture: the achievement of cognition-enabled robot manipulation in everyday situations, carrying out goal-directed tasks that need only be vaguely defined using underdetermined robot action plans specified in abstract terms. The vagueness is resolved at runtime by reasoning: querying knowledge bases and combining the resultant knowledge with information about the current state of the robot's environment acquired through perception, inferring the concrete actions that need to be performed to achieve the goal, and adapting them at runtime, as necessary. For example, figure 10.1 shows a PR2 robot setting a table during a demonstration of CRAM-based robot manipulation at the Everyday Activity Science and Engineering interdisciplinary research center (<https://ease-crc.org/>).

CRAM—see figure 10.2—comprises five core elements: 1) the CRAM Plan Language (CPL) executive; 2) a suite of knowledge bases and associated reasoning mechanisms, collectively referred to as KnowRob2 (Betz et al. 2018); 3) a perception executive; 4) an action executive; and 5) a metacognitive reasoning system. Several publications document



Figure 10.1

A PR2 robot setting a table during a demonstration of cognition-enabled robot manipulation using the CRAM. *Source:* Courtesy of the EASE interdisciplinary research center at the University of Bremen, Germany.

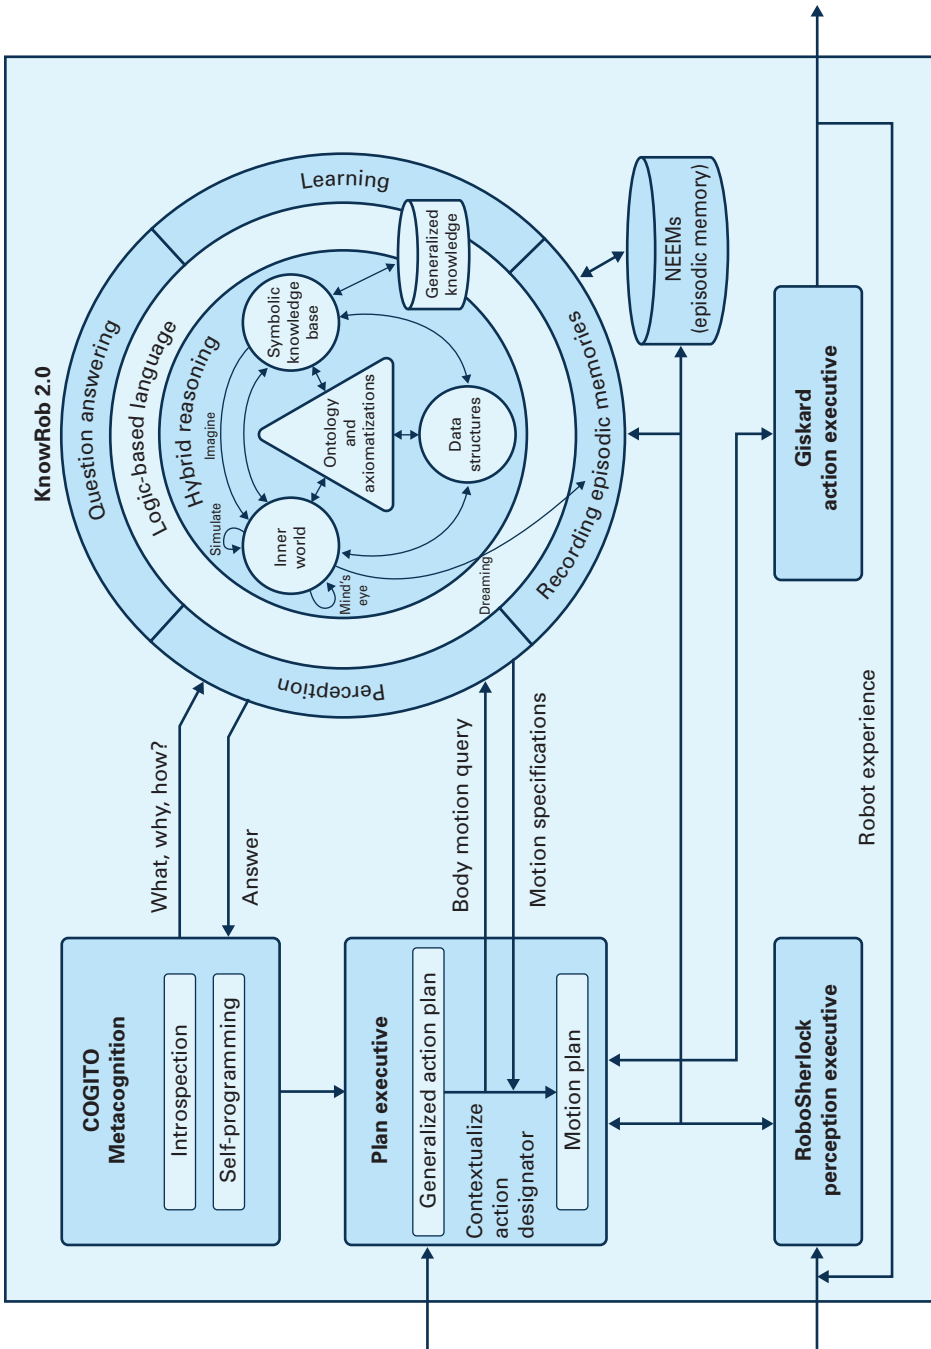


Figure 10.2
The CRAM cognitive architecture. *Source:* Courtesy of the EASE interdisciplinary research center at the University of Bremen, Germany.

the development of CRAM over the past ten years, a small subset of which includes Winkler et al. (2012), Tenorth and Beetz (2013), Beetz et al. (2015), and Kunze and Beetz (2017).

The CRAM Plan Language (CPL) executive is an extension of the Lisp programming language. It represents all the key aspects of a plan as persistent first-class objects in first-order logic. Thus, CRAM can reason about its plans, even at runtime. This is particularly relevant in the metacognition system. Plans specify how the robot should respond to sensory events, changes in belief states, and detected failures in plans. All these aspects of a plan can be queried, inspected, and reasoned about. A plan comprises a set of abstract plan designators for actions, objects, locations, and motions—that is, elementary movements. Designators are effectively placeholders and require runtime resolution based on the current context of the task action. Designator resolution is accomplished either by querying a priori knowledge embedded in the plan, by querying knowledge in the KnowRob2 knowledge base, or by accessing sensorimotor data through the perception executive. All plans have a similar generic structure, as shown below. The terms prefixed with a question mark are resolved at runtime based on the current state of the robot and the environment.

```
(par
  (perform
    (desig: an action
      (type picking-up)
      (arm ?grasping-arm)
      (grasp left-side)
      (object ?perceived-object ))
    . . .
  )
```

The KnowRob2 knowledge base is a knowledge representation and reasoning framework for robotic agents (cf. chapter 21). It is implemented in Prolog, and it is exposed as a conventional first-order time interval logic knowledge base. However, many logic expressions are constructed on demand from sensorimotor data computed in real time. It provides the background common sense intuitive-physics knowledge required by the CPL executive to implement its goal-directed underdetermined task plans—for example, how to grasp an object, depending on the object’s shape, weight, softness, and other properties; how it must be held while moving it—for example, upright to avoid spilling its contents; and where the object is normally located. Some knowledge is specified a priori, some is derived from experience, and some is the result of the simulated execution of candidate actions using a high-fidelity virtual reality physics engine simulator. All knowledge is represented by a first-order time interval logic expression and reasoned about, as needed.

KnowRob2 comprises five core elements embedded in a hybrid (i.e., multiformalism) reasoning shell, exposed through a logic-based language layer to an interface shell that provides perception, question answering, experience acquisition, and knowledge learning. The five elements are 1) a central set of knowledge ontologies and axiomatizations; 2) an episodic memory knowledge base encapsulating the robot’s experiences, represented in both subsymbolic form and in generalized symbolic form; 3) an inner-world knowledge base and virtual reality physics engine simulator; 4) a logic knowledge base with abstracted symbolic

sensor and action data, logical axioms, and inference rules; and 5) a virtual knowledge base comprising a set of data structures for parameterized motion control and path planning.

The knowledge ontologies and associated axiomatizations provide structured representation of the knowledge about the robot and its environment. There is a core ontology and additional special-purpose, application-specific ontologies. The core ontology defines the robot configuration, object configurations, robot actions, tasks, activities and behaviors, environment configuration, and situational context. The axioms identify roles that objects can play—for example, a mug is a cylindrical vessel, with a handle, that can be used as a receptacle from which its contents can be drunk, mixed, or poured.

One of the main distinguishing aspects of KnowRob2 is its focus on episodic memory. This is an autobiographical memory of the robot's experience as it had carried out tasks in the past. These are organized as NEEMS—narrative-enabled episodic memories—a concept introduced by the KnowRob2 designers. A NEEM comprises an experience and a narrative. The experience is a detailed low-level recording of a certain episode, such as records of poses and percepts based on exteroceptive and proprioceptive sensory data. It also includes control signals. This is unusual because motor aspects of memory are normally stored in procedural memory. Thus, CRAM episodic memory, and NEEMS in particular, generalizes the concept of an episode to include procedural elements. The narrative is an abstract symbolic description of the tasks, the context, the intended goals, and the observed effects (Beetz et al. 2018). KnowRob2 episodic memory, in representing procedural knowledge as declarative knowledge, allows it to be reasoned about. The episodic knowledge base provides the basis for answers to queries such as what actions were performed by the robot, when it performed them, how it performed them, why they were performed, whether or not they were successful, what the robot perceived while performing them, and what the robot believed when it performed them. The extraction of the generalized symbolic knowledge from NEEMS is facilitated by an interface to the Weka machine-learning framework (Holmes, Donkin, and Witten 1994).

The inner-world knowledge base facilitates geometric reasoning using a high-quality virtual reality system and physics engine. This allows KnowRob2 to simulate the outcome of candidate action and to establish the feasibility of that action. It provides symbolic names and properties for each entity, and it can infer background knowledge—for example, where an object is stored. The inner-world knowledge base serves two roles: as a representation of the belief state of the robot about itself and the world and as a reasoning mechanism for determining the outcome of candidate actions. Thus, it encapsulates two types of knowledge: current beliefs about robots and the world and the projected internal simulation of future states. It also acts as a learning mechanism, generating episodic memories off-line, effectively dreaming while physically inactive, and running simulations of activities with varying control parameters. These are recorded and transferred to the episodic-memory knowledge base.

The logic knowledge base provides information about the entities in the robot's environment, including objects, object parts, object articulation models, and environments composed of objects, actions, and events. It uses an entity description language that allows partial descriptions of entities in terms of both symbolic and subsymbolic properties.

The virtual knowledge base provides computable predicates that facilitate the integration of nonsymbolic data into the reasoning process, allowing symbolic queries of non-

symbolic data. This allows run-time sensorimotor states to be integrated into the knowledge base at run-time and to be used in reasoning in the same way as symbolic knowledge.

KnowRob2 provides a logic-based language interface that allows the hybrid reasoning shell to be exposed as a purely symbolic knowledge base even though internally it uses multiple symbolic and subsymbolic representations and reasoning formalisms. In this way, KnowRob2 can be treated by the CPL executive (and other systems through its OpenEASE interface; Beetz et al. 2015) as a symbolic, object-oriented query system in which entities can be retrieved by providing partial descriptions of them using the entity predicate. This allows KnowRob2 to appear as a “Siri for robots” (Beetz 2018)—that is, as a query and response oracle. Consequently, during task execution there is an ongoing dialogue between the CPL executive and KnowRob2, in which the CPL executive presents a series of under-determined queries, and KnowRob2 provides the corresponding responses, allowing the CPL executive to carry out the task using the action executive.

The action executive controls the robot by mapping parameterized actions (as requested by the CPL executive) to adaptive trajectories in real time.

Sensory information is available to the CPL executive either directly from the perception executive or indirectly through KnowRob2 by means of the virtual knowledge base and the associated computable predicates.

The metacognition subsystem allows CRAM to reason about plans and exploit transformational learning and planning to improve them in two complementary ways: by specialization using pragmatic everyday activity manifolds (PEAMs) and by generalization through metacognitive induction. This is possible because, as we noted above, the plans themselves are represented as first-class objects in first-order logic. PEAMs capture the subspace of motions necessary to carry out an action successfully by exploiting the constraints that knowledge of everyday activities and the environment bring to bear, rendering tractable by specialization the solution of problems that in their full generality are intractable. Generalization through metacognitive induction complements the PEAM solution strategy by exploring patterns among actions plans, seeking ways to transform them either by carrying out the action in a more efficient and effective manner or by accomplishing the outcome of the action in a different way.

10.6.2 ISAC

ISAC—intelligent soft arm control—is a hybrid cognitive architecture for an upper-torso humanoid robot also called ISAC (Kawamura et al. 2008). From a software engineering perspective, ISAC is constructed from an integrated collection of software agents and associated memories. Agents encapsulate all aspects of a component of the architecture, operate asynchronously (i.e., without a shared clock to keep the processing of all agents locked in step with each other), and communicate with each other by passing messages.

As shown in figure 10.3, the multiagent ISAC cognitive architecture comprises Activator Agents for motion control, Perceptual Agents, and a First-Order Response Agent (FRA) to effect reactive perception-action control. It has three memory systems: short-term memory (STM), long-term memory (LTM), and a working memory system (WMS).

STM has a robot-centered spatiotemporal memory of the perceptual events currently being experienced. This is called a Sensory EgoSphere (SES), and it is a discrete representation of what is happening around the robot, represented by a geodesic sphere indexed

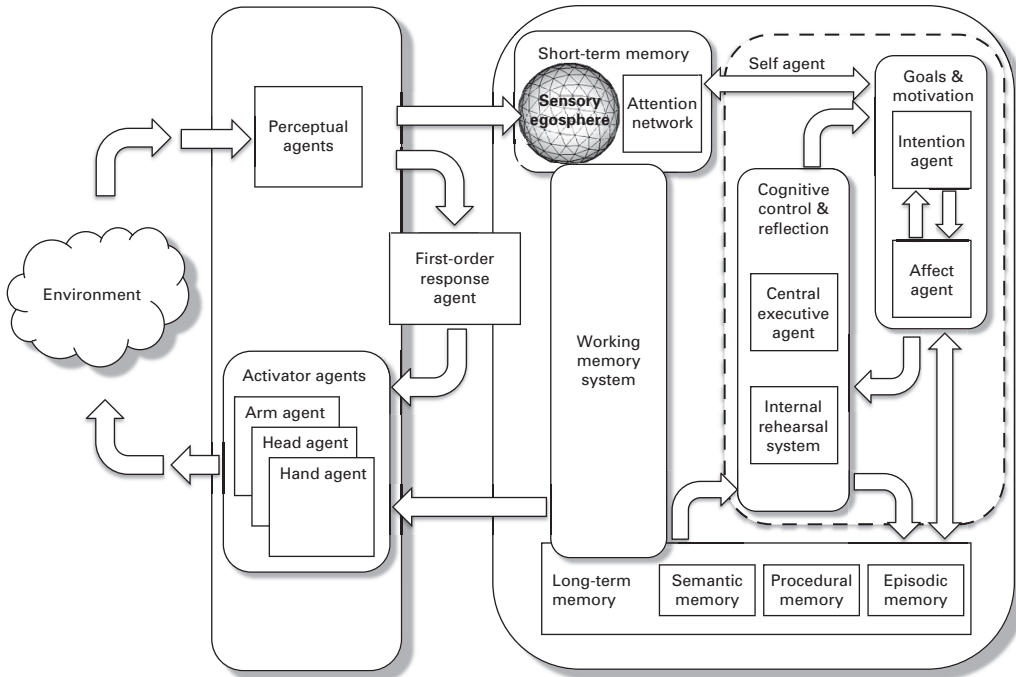


Figure 10.3
The ISAC cognitive architecture.

by two angles: horizontal (azimuth) and vertical (elevation). STM also has an Attention Network that determines the perceptual events that are most relevant and then directs the robot's attention to them.

LTM stores information about the robot's learned skills and past experiences. LTM is made up of semantic, episodic, and procedural memory. Together, the semantic memory and episodic memory make up the robot's declarative memory of the facts it knows. On the other hand, procedural memory stores representations of the motions the robot can perform.

ISAC's episodic memory abstracts past experiences and creates links or associations between them. It has multiple layers. At the lowest level, an episodic experience contains information about the external situation (i.e., task-relevant percepts from the SES), goals, emotions (in this case, internal evaluation of the perceived situation), actions, and outcomes that arise from actions and the valuations of these outcomes (e.g., how close they are to the desired goal state and any reward received as a result). Episodes are connected by links that encapsulate behaviors: transitions from one episode to another. Higher layers abstract away specific details and create links based on the transitions at lower levels. This multilayered approach allows for the efficient matching and retrieval of memories.

WMS, inspired by neuroscience models of brain function, temporarily stores information that is related to the task currently being executed. It forms a type of cache memory for STM, and the information it stores, called chunks, encapsulates expectations of future reward that are learned using a neural network.

Cognitive behavior is the responsibility of a Central Executive Agent (CEA) and an Internal Rehearsal System (IRS), a system that simulates the effects of possible actions. Together with a Goals and Motivation subsystem comprising an Intention Agent and an

Affect Agent, the CEA and the IRS form a compound agent called the Self Agent that, along with the FRA, makes decisions and acts according to the current situation and ISAC's internal states. The CEA is responsible for cognitive control, invoking the skills required to perform some given task on the basis of the current focus of attention and past experiences. The goals are provided by the Intention Agent. Decision-making is modulated by the Affect Agent.

ISAC works in the following way. Normally, the FRA produces reactive responses to sensory triggers. However, it is also responsible for executing tasks. When a task is assigned by a human, the FRA retrieves the skill from procedural memory in the LTM that corresponds to the skill described in the task information. It then places it in the WMS as chunks along with the current percept. The Activator Agent then executes it, suspending execution whenever a reactive response is required. If the FRA finds no matching skill for the task, the CEA takes over, recalling from episodic memory past experiences and behaviors that contain information similar to the current task. One behavior-percept pair is selected, based on the current percept in the SES, its relevance, and the likelihood of successful execution as determined by internal simulation in the IRS. This is then placed in working memory, and the Activator Agent executes the action.

10.7 Future Prospects

The design and implementation of a cognitive architecture is a daunting undertaking. This is evident when you consider that contemporary cognitive architectures such as Soar (Laird, Newell, and Rosenbloom 1987; Laird 2009, 2012), ACT-R (Anderson et al. 2004; Anderson 1996), CLARION (Sun 2007, 2016), and CRAM (Beetz et al. 2010; Mösenlechner 2016) have taken ten to twenty years or more to develop and are all still being developed further.¹ In an effort to consolidate cognitive architecture research, the cognitive science community has launched an exercise to identify the key design features shared by the most prominent cognitive architectures, with the goal of creating a common model of cognition (Laird, Lebiere, and Rosenbloom 2017) and promoting more cohesive development and achieving greater progress. In any case, progress will depend on the thorough evaluation of cognitive architectures in diverse, challenging, realistic environments (Kotseruba and Tsotsos 2020) consistent with human-level intelligence, such as the CRAM cognitive architecture targets in everyday activity science and engineering (EASE).

There is a need for more realistic perceptual capabilities that can operate in adverse conditions with noise and uncertainty, using context to improve performance. Almost half the cognitive architectures surveyed by Kotseruba and Tsotsos do not implement any visual perception or other sensory modalities. For example, audition, touch, and olfaction are typically addressed in a trivial manner (Kotseruba and Tsotsos 2020).

Cognitive architectures also need to facilitate more natural communication with humans to infer their intentions and emotional states; engage in adaptive, personalized interaction; read body language, such as gestures and facial expressions; engage in natural turn-taking; and facilitate human-robot joint action. Examples of cognitive architectures that focus on these aspects of cognitive human-robot interaction include Lemaignan et al. (2017); Sandini et al. (2018); Tanevska et al. (2019).

Computational models of episodic memory have not received significant attention, especially for lifelong learning, despite the fact that its existence and importance has been

widely recognized (Kotseruba and Tsotsos 2020). Notable exceptions include the CRAM cognitive architecture (Beetz, Mösenlechner, and Tenorth 2010; Mösenlechner 2016) and the iCub neural framework for episodic memory (Mohan, Sandini, and Morasso 2014).

Deep learning (Schmidhuber 2015; Goodfellow, Bengio, and Courville 2016) has not yet made a significant impact on cognitive architectures (Kotseruba and Tsotsos 2020). This will almost certainly change, giving rise to new architectural requirements—for example, deep developmental robotics architectures (Sigaud and Droniou 2016) and a reconciliation of deep learning with symbolic artificial intelligence (Garnelo and Shanahan 2019). One of the main advantages of deep learning is its ability to produce end-to-end systems—that is, systems that map directly from an input space to an output space, such as pixels-to-classes in computer vision. In robotics, the situation is different: end-to-end systems must map from pixels (and other sensory stimuli) to torques in a dynamic interactive environment. Supervised deep learning based on static data sets is not viable in these circumstances. However, deep reinforcement learning (Arulkumaran et al. 2017; Li 2018) is capable of learning end-to-end robot control or action policies. This form of learning is typically implemented using simulators and may not be feasible on physical robots. Sünderhauf et al. (2018) estimate that it would take fifty-three days to accomplish a deep reinforcement learning exercise that currently takes twenty-four hours using simulation. They suggest that there is also a reality gap between simulation and the real world that limits the usefulness of simulation-based deep reinforcement learning, and they discuss a solution based on transfer learning, initially learning in the simulated environment, freezing the network weights, and then continuing the learning with the physical robot. On the other hand, results using photorealistic simulations to support reasoning in cognition-enabled robots (Beetz et al. 2018; Mania and Beetz 2019) suggest that the reality gap may not be significant and that the simulation approach may be plausible.

10.8 Conclusion

A cognitive architecture captures both the abstract conceptual form and the details of functional operation, focusing on inner cohesion and self-contained completeness. This means that all of the mechanisms required for cognition fall under the compass of a cognitive architecture, including perception, attention, action, control, learning, reasoning, memory, adaptivity, and anticipation. Thus, cognition, as a process, and a cognitive architecture, as a framework, embrace all of the elements required for effective action. A cognitive architecture specifies the system components and the way these components are dynamically related as a whole. It provides both an abstract model of cognitive behavior and a sufficient basis for a software instantiation of that model (Lieto et al. 2018). Despite the magnitude of the task, the design and implementation of an appropriate cognitive architecture remains an indispensable step in the creation of a cognitive robot.

Additional Reading and Resources

- To delve more deeply into the field of cognitive architectures, you might begin by reading the review by Kotseruba and Tsotsos (2020) and referring to the companion

website: http://jtl.lassonde.yorku.ca/project/cognitive_architectures_survey/index.html. The review does not focus specifically on robot cognitive architectures but provides a contemporary and comprehensive overview of the field, nonetheless.

- Appendix A of Vernon, von Hofsten, and Fadiga (2011), summarizing the operation of twenty cognitive architectures: Vernon, D., C. von Hofsten, and L. Fadiga. 2011. *A Roadmap for Cognitive Development in Humanoid Robots*. In Vol. 11, *Cognitive Systems Monographs*. Berlin: Springer. http://www.vernon.eu/COSMOS_CAs.pdf.
- The Introduction to Cognitive Robotics course (www.cognitiverobotics.net) has several lectures devoted to cognitive architectures, in general, and to the CRAM cognitive architecture summarized in section 10.6.1, in particular, expanding on the material in the online CRAM tutorials: <http://cram-system.org/tutorials>.
- Software is available online for, for example, the CRAM cognitive architecture: <http://cram-system.org>; the openEASE software components for cognition-enabled control of robotic agents: <https://ease-crc.org/open-ease/>; and the iCub cognitive robot platform: <http://www.icub.org>.
- Instructions on how to access, download, and install the CRAM software is included in the Introduction to Cognitive Robotics course and on the CRAM website (<http://cram-system.org/installation>), along with practical exercises to help you get started.
- For other software resources, refer to the “Resources” page on the IEEE Technical Committee for Cognitive Robotics site: <http://www.ieee-coro.org>.

Note

1. The average age of cognitive architecture projects in the survey by Kotseruba and Tsotsos (2020) is approximately fifteen years.

References

- Anderson, John R. 1996. “ACT: A Simple Theory of Complex Cognition.” *American Psychologist* 51 (4): 355.
- Anderson, John R., Daniel Bothell, Michael D. Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. 2004. “An Integrated Theory of the Mind.” *Psychological Review* 111 (4): 1036.
- Arulkumaran, Kai, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. “Deep Reinforcement Learning: A Brief Survey.” *IEEE Signal Processing Magazine* 34 (6): 26–38.
- Atance, Cristina M., and Daniela K. O’Neill. 2001. “Episodic Future Thinking.” *Trends in Cognitive Sciences* 5 (12): 533–539.
- Beetz, Michael. 2018. Personal Communication.
- Beetz, Michael, Daniel Beßler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoğlu, and Georg Bartels. 2018. “KnowRob 2.0—a 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents.” In *2018 IEEE International Conference on Robotics and Automation*, 512–519. New York: IEEE.
- Beetz, Michael, Dominik Jain, Lorenz Mösenlechner, and Moritz Tenorth. 2010. “Towards Performing Everyday Manipulation Activities.” *Robotics and Autonomous Systems* 58 (9): 1085–1095.
- Beetz, Michael, Lorenz Mösenlechner, and Moritz Tenorth. 2010. “CRAM—a Cognitive Robot Abstract Machine for Everyday Manipulation in Human Environments.” In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1012–1017. New York: IEEE.
- Beetz, Michael, Moritz Tenorth, and Jan Winkler. 2015. “Open-EASE—a Knowledge Processing Service for Robots and Robotics/AI Researchers.” In *2015 IEEE International Conference on Robotics and Automation*, 1983–1990. New York: IEEE.

- Borji, Ali, Dicky N. Sihite, and Laurent Itti. 2013. "What/Where to Look Next? Modeling Top-Down Visual Attention in Complex Interactive Environments." *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 44 (5): 523–538.
- Cangelosi, Angelo, and Matthew Schlesinger. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Christensen, Wayne D., and Clifford Alan Hooker. 2000. "An Interactivist-Constructivist Approach to Intelligence: Self-Directed Anticipative Learning." *Philosophical Psychology* 13 (1): 5–45.
- Demiris, Yiannis, and Bassam Khadhouri. 2006. "Hierarchical Attentive Multiple Models for Execution and Recognition of Actions." *Robotics and Autonomous Systems* 54 (5): 361–369.
- Doya, Kenji. 1999. "What Are the Computations of the Cerebellum, the Basal Ganglia and the Cerebral Cortex?" *Neural Networks* 12 (7–8): 961–974.
- Duch, Włodzisław, Richard Jayadi Oentaryo, and Michel Pasquier. 2008. "Cognitive Architectures: Where Do We Go from Here?" In *Artificial General Intelligence 2008, Proceedings of the First AGI Conference, AGI 2008, March 1–3, 2008, University of Memphis, Memphis, TN, USA*, 122–136. *Frontiers in Artificial Intelligence and Applications* 171. Amsterdam: IOS Press.
- EASE. n.d. Everyday Activity Science and Engineering. Accessed August 21, 2021. <https://ease-crc.org/>.
- Edelman, Gerald M. 2006. *Second Nature: Brain Science and Human Knowledge*. New Haven, CT: Yale University Press.
- Eliasmith, Chris, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis Dewolf, Yichuan Tang, and Daniel Rasmussen. 2012. "A Large-Scale Model of the Functioning Brain." *Science* 338 (6111): 1202–1205.
- Flanagan, J. Randall, and Roland S. Johansson. 2003. "Action Plans Used in Action Observation." *Nature* 424 (6950): 769–771.
- Flanagan, J. Randall, Gerben Rotman, Andreas F. Reichelt, and Roland S. Johansson. 2013. "The Role of Observers' Gaze Behavior When Watching Object Manipulation Tasks: Predicting and Evaluating the Consequences of Action." *Philosophical Transactions of the Royal Society B: Biological Sciences* 368 (1628): 20130063.
- Franklin, Stan, Tamas Madl, Sidney D'Mello, and Javier Snaider. 2014. "LIDA: A Systems-Level Architecture for Cognition, Emotion, and Learning." *IEEE Transactions on Autonomous Mental Development* 6 (1): 19–41.
- Franklin, Stan, Uma Ramamurthy, Sidney K. D'Mello, Lee Mccauley, Aregahegn Negatu, Rodrigo L. Silva, and Vivek Datla. 2007. "LIDA: A Computational Model of Global Workspace Theory and Developmental Learning." In *AAAI Fall Symposium on AI and Consciousness: Theoretical Foundations*, 61–66. Menlo Park, CA: AAAI Press.
- Froese, Tom, Nathaniel Virgo, and Eduardo Izquierdo. 2007. "Autonomy: A Review and a Reappraisal." In Vol. 4648, *Proceedings of the 9th European Conference on Artificial Life: Advances in Artificial Life*, edited by F. Almeida E. Costa, L. Rocha, E. Costa, I. Harvey, and A. Coutinho, 455–465. Berlin: Springer.
- Garnelo, Marta, and Murray Shanahan. 2019. "Reconciling Deep Learning with Symbolic Artificial Intelligence: Representing Objects and Relations." *Current Opinion in Behavioral Sciences* 29:17–23.
- Gilbert, Daniel T., and Timothy D. Wilson. 2007. "Prospection: Experiencing the Future." *Science* 317 (5843): 1351–1354.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.
- Hesslow, Germund. 2002. "Conscious Thought as Simulation of Behavior and Perception." *Trends in Cognitive Sciences* 6 (6): 242–247.
- Hesslow, Germund. 2012. "The Current Status of the Simulation Theory of Cognition." *Brain Research* 1428:71–79.
- Holmes, Geoffrey, Andrew Donkin, and Ian H. Witten. 1994. "Weka: A Machine Learning Workbench." In *Proceedings of ANZILS'94—Australian New Zealand Intelligent Information Systems Conference*, 357–361. New York: IEEE.
- Kawamura, Kazuhiko, Stephen M. Gordon, Palis Ratanaswasd, Erdem Erdemir, and Joseph F. Hall. 2008. "Implementation of Cognitive Control for a Humanoid Robot." *International Journal of Humanoid Robotics* 5 (4): 547–586.
- Kieras, Davis E., and Davis E. Meyer. 1997. "An Overview of the EPIC Architecture for Cognition and Performance with Application to Human-Computer Interaction." *Human-Computer Interaction* 12 (4): 391–438.
- Kotseruba, Iuliia, and John K. Tsotsos. 2020. "40 Years of Cognitive Architectures: Core Cognitive Abilities and Practical Applications." *Artificial Intelligence Review* 53 (1): 17–94.
- Kunze, Lars, and Michael Beetz. 2017. "Envisioning the Qualitative Effects of Robot Manipulation Actions Using Simulation-Based Projections." *Artificial Intelligence* 247:352–380.
- Laird, John E. 2009. "Toward Cognitive Robotics." In Vol. 7332, *Proceedings of the SPIE—Unmanned Systems Technology XI*. P. 73320Z. Bellingham, WA: International Society for Optics and Photonics.
- Laird, John E. 2012. *The Soar Cognitive Architecture*. Cambridge, MA: MIT Press.

- Laird, John E., Christian Lebiere, and Paul S. Rosenbloom. 2017. "A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics." *AI Magazine* 38 (4): 13–26.
- Laird, John E., Allen Newell, and Paul S. Rosenbloom. 1987. "Soar: An Architecture for General Intelligence." *Artificial Intelligence* 33:1–64.
- Langley, Pat. 2006. "Cognitive Architectures and General Intelligent Systems." *AI Magazine* 27 (2): 33–33.
- Langley, Pat, and Dongkyu Choi. 1999 [2006]. "A Unified Cognitive Architecture for Physical Agents." In *Proceedings of the National Conference on Artificial Intelligence*, Vol. 21, no. 2, 1469. Menlo Park, CA: AAAI Press; Cambridge, MA: MIT Press.
- Langley, Pat, John E. Laird, and Seth Rogers. 2009. "Cognitive Architectures: Research Issues and Challenges." *Cognitive Systems Research* 10 (2): 141–160.
- Lemaignan, Séverin, Mathieu Warnier, E. Akin Sisbot, Aurélie Clodic, and Rachid Alami. 2017. "Artificial Cognition for Social Human–Robot Interaction: An Implementation." *Artificial Intelligence* 247:45–69.
- Li, Yuxi. 2018. "Deep Reinforcement Learning." ArXiv preprint: 1801.06339v1.
- Lieto, Antonio, Mehul Bhatt, Alessandro Oltramari, and David Vernon. 2018. "The Role of Cognitive Architectures in General Artificial Intelligence." *Cognitive Systems Research* 48:1–3.
- Lindblom, Jessica. 2015. *Embodied Social Cognition*. Vol. 26. Berlin: Springer.
- Mania, Patrick, and Michael Beetz. 2019. "A Framework for Self-Training Perceptual Agents in Simulated Photorealistic Environments." In *2019 International Conference on Robotics and Automation*, 4396–4402. New York: IEEE.
- Merrick, Kathryn Elizabeth. 2010. "A Comparative Study of Value Systems for Self-Motivated Exploration and Learning by Robots." *IEEE Transactions on Autonomous Mental Development* 2 (2): 119–131.
- Mohan, Vishwanathan, Ajaz Bhat, and Pietro Morasso. 2019. "Muscleless Motor Synergies and Actions without Movements: From Motor Neuroscience to Cognitive Robotics." *Physics of Life Reviews* 30:89–111.
- Mohan, Vishwanathan, Giulio Sandini, and Pietro Morasso. 2014. "A Neural Framework for Organization and Flexible Utilization of Episodic Memory in Cumulatively Learning Baby Humanoids." *Neural Computation* 26 (12): 2692–2734.
- Mösenlechner, Lorenz. 2016. "The Cognitive Robot Abstract Machine: A Framework for Cognitive Robotics." PhD thesis, Technical University of Munich.
- Newell, Allen. 1982. "The Knowledge Level." *Artificial Intelligence* 18 (1): 87–127.
- Newell, Allen. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Newell, Allen. 1992. "Précis of Unified Theories of Cognition." *Behavioral and Brain Sciences* 15:425–492.
- Oudeyer, Pierre-Yves, Frédéric Kaplan, and Verena V. Hafner. 2007. "Intrinsic Motivation Systems for Autonomous Mental Development." *IEEE Transactions on Evolutionary Computation* 11 (2): 265–286.
- Piaget, Jean. 1954. *The Construction of Reality in the Child*. New York: Basic Books.
- Piccinini, Gualtiero. 2010. "The Mind as Neural Software? Understanding Functionalism, Computationalism, and Computational Functionalism." *Philosophy and Phenomenological Research* 81 (2): 269–311.
- Ritter, Frank E., and Richard M. Young. 2001. "Embodied Models as Simulated Users: Introduction to This Special Issue on Using Cognitive Models to Improve Interface Design." *International Journal of Human-Computer Studies* 55 (1): 1–14.
- Rosenbloom, Paul S., Abram Demski, and Volkan Ustun. 2016. "The Sigma Cognitive Architecture and System: Towards Functionally Elegant Grand Unification." *Journal of Artificial General Intelligence* 7 (1): 1–103.
- Ryle, G. 1949. *The Concept of Mind*. London: Hutchinson's University Library.
- Samsonovich, Alexei V. 2010. "Toward a Unified Catalog of Implemented Cognitive Architectures." In *Proceedings of the Conference on Biologically Inspired Cognitive Architectures*, 195–244. Amsterdam: IOS Press.
- Sandini, Giulio, Vishwanathan Mohan, Alessandra Sciutti, and Pietro Morasso. 2018. "Social Cognition for Human-Robot Symbiosis—Challenges and Building Blocks." *Frontiers in Neurobotics* 12:34.
- Schacter, Daniel L., Donna Rose Addis, and Randy L. Buckner. 2008. "Episodic Simulation of Future Events: Concepts, Data, and Applications." *Annals of the New York Academy of Sciences* 1124:39–60.
- Schacter, Daniel L., Donna Rose Addis, Demis Hassabis, Victoria C. Martin, R. Nathan Spreng, and Karl K. Szpunar. 2012. "The Future of Memory: Remembering, Imagining, and the Brain." *Neuron* 76 (4): 677–694.
- Schmidhuber, Jürgen. 2015. "Deep Learning in Neural Networks: An Overview." *Neural Networks* 61:85–117.
- Seligman, Martin E. P., Peter Railton, Roy F. Baumeister, and Chandra Sripada. 2013. "Navigating into the Future or Driven by the Past." *Perspectives on Psychological Science* 8 (2): 119–141.

- Shanahan, Murray. 2006. "A Cognitive Architecture That Combines Internal Simulation with a Global Work-space." *Consciousness and Cognition* 15 (2): 433–449.
- Sigaud, Olivier, and Alain Droniou. 2016. "Towards Deep Developmental Learning." *IEEE Transactions on Cognitive and Developmental Systems* 8 (2): 99–114.
- Sterling, Peter. 2012. "Allostasis: A Model of Predictive Regulation." *Physiology and Behavior* 106 (1): 5–15.
- Stewart, John, Olivier Gapenne, and Ezequiel A. Di Paolo. 2010. *Enaction: Toward a New Paradigm for Cognitive Science*. Cambridge, MA: MIT Press.
- Sun, Ron. 2004. "Desiderata for Cognitive Architectures." *Philosophical Psychology* 17 (3): 341–373.
- Sun, Ron. 2007. "The Importance of Cognitive Architectures: An Analysis Based on CLARION." *Journal of Experimental and Theoretical Artificial Intelligence* 19 (2): 159–193.
- Sun, Ron. 2016. *Anatomy of the Mind: Exploring Psychological Mechanisms and Processes with the Clarion Cognitive Architecture*. Oxford: Oxford University Press.
- Sünderhauf, Niko, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, et al. 2018. "The Limits and Potentials of Deep Learning for Robotics." *International Journal of Robotics Research* 37 (4–5): 405–420.
- Svensson, Henrik, Jessica Lindblom, and Tom Ziemke. 2007. "Making Sense of Embodied Cognition: Simulation Theories of Shared Neural Mechanisms for Sensorimotor and Cognitive Processes." In *Body, Language and Mind, Volume 1: Embodiment*, edited by T. Ziemke, J. Zlatev, and R. M. Frank, 241–269. Berlin: Mouton De Gruyter.
- Tanevska, Ana, Francesco Rea, Giulio Sandini, Lola Cañamero, and Alessandra Sciutti. 2019. "A Cognitive Architecture for Socially Adaptable Robots." In *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics*, 195–200. New York: IEEE.
- Tenorth, Moritz, and Michael Beetz. 2013. "KnowRob: A Knowledge Processing Infrastructure for Cognition-Enabled Robots." *International Journal of Robotics Research* 32 (5): 566–590.
- Thórisson, Kristinn, and Helgi Helgasson. 2012. "Cognitive Architectures and Autonomy: A Comparative Review." *Journal of Artificial General Intelligence* 3 (2): 1–30.
- Trafton, J. Gregory, Nicholas L. Cassimatis, Magdalena D. Bugajska, Derek P. Brock, Farilee E. Mintz, and Alan C. Schultz. 2005. "Enabling Effective Human-Robot Interaction Using Perspective-Taking in Robots." *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 35 (4): 460–470.
- Tsotsos, John K. 2011. *A Computational Perspective on Visual Attention*. Cambridge, MA: MIT Press.
- Tulving, Endel. 1972. "Episodic and Semantic Memory." *Organization of Memory* 1:381–403.
- Tulving, Endel. 1984. "Précis of Elements of Episodic Memory." *Behavioral and Brain Sciences* 7 (2): 223–238.
- Vernon, David. 2010. "Enaction as a Conceptual Framework for Development in Cognitive Robotics." *Paladyn Journal of Behavioral Robotics* 1 (2): 89–98.
- Vernon, David. 2014. *Artificial Cognitive Systems—a Primer*. Cambridge, MA: MIT Press.
- Vernon, David, Giorgio Metta, and Giulio Sandini. 2007a. "The iCub Cognitive Architecture: Interactive Development in a Humanoid Robot." In *2007 IEEE 6th International Conference on Development and Learning*, 122–127. New York: IEEE.
- Vernon, David, Giorgio Metta, and Giulio Sandini. 2007b. "A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents." *IEEE Transactions on Evolutionary Computation* 11 (2): 151–180.
- Vernon, David, Claes von Hofsten, and Luciano Fadiga. 2011. *A Roadmap for Cognitive Development in Humanoid Robots*. Vol. 11. Berlin: Springer Science and Business Media.
- Vernon, David, Claes von Hofsten, and Luciano Fadiga. 2016. "Desiderata for Developmental Cognitive Architectures." *Biologically Inspired Cognitive Architectures* 18:116–127.
- von Hofsten, Claes. 2009. "Action, the Foundation for Cognitive Development." *Scandinavian Journal of Psychology* 50:617–623.
- Vygotsky, Lev S. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press.
- Wiener, Norbert. 1948. *Cybernetics: Or the Control and Communication in the Animal and the Machine*. New York: John Wiley and Sons.
- Winkler, Jan, Georg Bartels, Lorenz Mösenlechner, and Michael Beetz. 2012. "Knowledge Enabled High-Level Task Abstraction and Execution." In *First Annual Conference on Advances in Cognitive Systems* 2 (1): 131–148.

11 Embodiment in Cognitive Science and Robotics

Tom Ziemke

11.1 Introduction

Writing a book chapter on the notion of *embodiment* in the cognitive sciences, or cognitive robotics more specifically, is not an easy task these days. Many researchers nowadays share the belief that, as M. Wilson (2002, 625) formulated it, “Cognitive processes are deeply rooted in the body’s interactions with the world,” and we can take that as a useful first approximation of the fundamental claim of *embodied cognition* research. That statement, however, means surprisingly many and surprisingly different things to different people.

Hence, explaining what embodiment is, in a single chapter, is difficult for several reasons: First, embodiment has been discussed in the cognitive sciences for several decades now. Early examples include Lakoff and Johnson’s (1980) work on the role of bodily metaphors in human cognition and language, as well as Maturana and Varela’s (1980) work on the biology of cognition. Moreover, many of these discussions have roots preceding cognitive science as a discipline by several more decades, such as the work of William James in the late nineteenth century. So there simply is a long history to cover. Second, embodied cognition has become a popular and more-or-less mainstream position in the last twenty years (e.g., Clark 1999; Damasio 1999; Gallagher 2005; Ziemke et al. 2006; Pfeifer and Bongard 2007; Chemero 2009; Shapiro 2010; Black 2014; Lindblom 2015). Some would go so far as to claim that “embodied cognition is sweeping the planet,” as it says in one of the endorsements on the back cover of Shapiro’s (2010) textbook on the topic. However, research on embodied cognition is not so well established and mainstream yet that research has converged sufficiently to establish clear boundaries and shared definitions of what is or is not embodied cognition (e.g., Wilson and Golonka 2013; Ziemke and Thill 2014; Ziemke 2016). Hence, there are many different—and in some cases also conflicting—perspectives to address. Last, but not least, the issue of embodiment is somewhat uniquely placed at the intersection of engineering, science, and philosophy, which means that embodiment simply has different significance to different, but overlapping, research communities.

The latter point should be relatively clear in the context of a book on cognitive robotics (see also section 1.2): On the one hand, there is the *engineering perspective* (with an emphasis

on the “robotics” in “cognitive robotics”) on how to equip robots with the required sensorimotor, cognitive, and communicative capacities for particular tasks. If you think of a typical example, such as a service robot helping elderly people at home, it is clear that the robot’s embodiment—in the sense of what the robot looks like, what sensors, actuators, and interactive modalities it has, and so on—plays a crucial role in determining what it can do and how people can interact with it. After all, robot lawnmowers and vacuum cleaners, for example, might be well suited for their specific purposes, but they are not exactly easy to talk to. On the other hand, there is the (cognitive-) *scientific perspective* (with an emphasis on the “cognitive” in “cognitive robotics”), according to which embodied (i.e., robotic) models that share some bodily and sensorimotor features with the organism they are supposed to model might be preferable to purely computational or mathematical models. If, for example, you are working on modeling how human language use is grounded in sensorimotor interaction, then it might make sense to use humanoid robot models that are, at least to some degree, similar in terms of bodily features and sensorimotor capacities to the people and processes you are trying to model. On the third hand (to use an intentionally confusing bodily metaphor), there is of course the *philosophical perspective*, according to which theories of embodied cognition and cognitive-robotics models offer novel approaches to age-old questions concerning the so-called mind-body problem.

To cover a broad range of perspectives on embodiment, the remainder of this chapter is structured as follows: section 11.2 asks some basic questions—such as what is a body, what do we mean by embodiment, and what do we mean by embodied cognition—and provides some preliminary answers in the form of basic distinctions that might be useful. Section 11.3 then addresses fundamental conceptions of embodied cognition in cognitive science and the philosophy of mind. Section 11.4 narrows the perspective to notions of embodiment in artificial intelligence (AI) research, where naturally some of the central questions are what would constitute an artificial body or embodiment capable of supporting artificial embodied intelligence, and how we should go about building such systems. Section 11.5 then addresses the role of embodiment in cognitive robotics more specifically and connects back to the above discussion of different perspectives on embodiment. Section 11.6, finally, provides a brief summary and some conclusions.

It should be noted that throughout this chapter, for the reasons mentioned above, we prioritize breadth—that is, provide a broad spectrum of perspectives on embodiment and its role in human (and robotic) cognition and refer the interested reader to the original literature for more in-depth discussions.

11.2 Notions of Embodiment and Embodied Cognition

What actually is a body? Or, more specifically, what constitutes the kind of body—or embodiment—that might be a necessary requirement for embodied cognition? Somewhat surprisingly, maybe, many discussions of embodied cognition actually pay relatively little attention to the nature and the role of the body involved. This might be natural in psychology or linguistics, which mainly deal with phenomena such as human cognition or language, where it is more or less obvious that discussions of embodiment are about the role that *human bodies* play in such phenomena. In AI and robotics, however, things are less

obvious. This raises questions, such as what kind of embodiment might be required for an artificial system that could deal with, for example, human language. In the realm of science fiction, you might have noticed that the *Star Wars* android C-3PO claims to be fluent in six million forms of communication and thus is presumably capable of dealing with many different species, although his embodiment is rather humanlike. In today's real world, on the other hand, many of us regularly encounter systems that appear rather disembodied, such as Google Translate, or devices with somewhat minimal and distinctly nonhuman physical "embodiments," such as Amazon Echo or Google Home, which nevertheless all seem to be able to deal with human language to some degree. Let us therefore have a quick look at some of the notions of what kind of body or embodiment is required for embodied cognition (following Ziemke 2001, 2003b).

Embodiment as Structural Coupling

Probably the broadest notion of embodiment is that systems are embodied if they are structurally coupled to their environment. The concept of structural coupling originally comes from Maturana and Varela's (1980, 1987) work on the biology of cognition, which will be discussed in further detail below. Quick and colleagues (1999, 340) tried to formalize this as follows in their minimal definition of embodiment: "A system X is embodied in an environment E if perturbatory channels exist between the two. That means, X is embodied in E if for every time t at which both X and E exist, some subset of E's possible states with respect to X have the capacity to perturb X's state, and some subset of X's possible states with respect to E have the capacity to perturb E's state." This definition of embodiment in itself does not make a distinction between cognitive and noncognitive systems, which is illustrated by Quick et al.'s (1999) example of a granite outcrop on the Antarctic tundra that is perturbed by the wind and in turn perturbs the flow of air. This would seem to include practically all physical objects, but it might be worth noting that it has been argued that structural coupling does not necessarily require a physical body. Franklin (1997, 500), for example, explicitly stated: "Software systems with no body in the usual physical sense can be intelligent. But they must be embodied in the situated sense of being autonomous agents structurally coupled with their environment."

Historical Embodiment

Several researchers have emphasized that cognitive systems are not only structurally coupled to their environment in the present. Their embodiment is in fact a result or reflection of a history of agent-environment interaction. According to Varela et al. (1991, 149), for example, "Knowledge depends on being in a world that is inseparable from our bodies, our language, and our social history—in short, from our embodiment." Ziemke (1999, 187) pointed out: "Natural embodiment is more than being-physical . . . it reflects/embodies the history of structural coupling and mutual specification between agent and environment in the course of which the body has been constructed." In a similar vein, Riegler (2002, 347) included an agent's adaptation to its environment over time in his definition of embodiment: "A system is embodied if it has gained competence within the environment in which it has developed."

Physical/Sensorimotor Embodiment

Many researchers in embodied (robotic) AI—to distinguish their approach from traditional AI—hold that, as Pfeifer and Scheier (1999, 649) formulated it, “intelligence cannot merely exist in the form of an abstract algorithm but requires a physical instantiation, a body.” This would seem to rule out software agents but could possibly still be considered to include the granite outcrop mentioned above. However, most researchers in embodied AI and robotics naturally adopt a more restrictive version of the notion of physical embodiment—which might be labeled *sensorimotor embodiment*—that is, the view that embodied systems need to be connected to their environment not just through physical forces but more specifically through their own sensors and actuators. This is also the essence of Brooks’s (1990) *physical grounding hypothesis*, according to which building an intelligent system requires having its representations grounded in the physical world, which in turn requires connecting it to the world via sensors and actuators.

At this point it might be worth pointing out that, although historical embodiment and physical/sensorimotor embodiment can be considered special cases of structural coupling, neither of these notions necessarily includes or excludes the other. Riegler (2002), for example, stated that his historical notion definition of embodiment “does not exclude domains other than the physical domain” and in particular that “computer programs may also become embodied” if they are the result of self-organization rather than conventional human design and programming.

Organismoid Embodiment

Another, yet more restrictive, notion of physical and sensorimotor embodiment would be that at least certain types of organism-like cognition might be limited to organism-like bodies—that is, physical bodies with particular structural features or sensorimotor capacities. A simple early example of this was a robot used by Lund et al. (1998) equipped with an auditory circuit and two microphones the same distance from each other as the two “ears” of the cricket whose phonotaxis it was supposed to model. The placement of the sensors/ears, in both robot and cricket, reduced the amount of internal signal processing required to respond selectively to certain sound frequencies. Note that in this case the bodies of the cricket and the wheeled robot were of course very different except for one crucial detail, namely the distance between the “ears.” More recent and complex examples of organismoid embodiment can be found in humanoid cognitive robotics, such as the work of Cangelosi and colleagues, in which humanlike embodiment is taken to be a key ingredient for robotic models of human language learning and use (e.g., Morse et al. 2015; cf. section 11.5).

Organismic Embodiment

The most restrictive notion of embodiment discussed in this section is that embodied cognition emerges from the interaction between organisms—that is, living bodies, and their environments. Maturana and Varela’s (1980, 1987) work on the biology of cognition, for example, suggests, in a nutshell, that cognition is what living systems do in interaction with their environment. In a similar vein, from a neuroscientific perspective Damasio (1998) criticized “the prevalent absence of a notion of organism in the sciences of mind and brain” as a problem, which he elaborated as follows: “It is not just that the mind remained linked to the brain

in a rather equivocal relationship, but that the brain remained consistently separated from the body and thus not part of the deeply interwoven mesh of body and brain that defines a complex living organism” (Damasio 1998, 84). Somatic theories of emotion constitute “a multi-tiered affectively embodied view of mind” (Panksepp 2005, 63), according to which emotion, cognition, and consciousness arise from multiple, nested levels of homeostatic (self-) regulation of bodily activity. This is, at least at this point in time, a clear difference between living systems and robotic bodies, which typically have no *own* needs or viability constraints (e.g., Bickhard 2009; Ziemke 2016) and therefore need to be equipped with artificial motivational systems (cf. chapter 13).

In addition to the above different notions of embodiment, let us also take a quick look at different views of embodied cognition. M. Wilson (2002) distinguished six such views from a psychological perspective, of which, however, only the last explicitly addresses the role of body:

“Cognition is situated” This claim is widely held in the literature on embodied cognition. M. Wilson (2002) herself distinguished between situated cognition, taking place “in the context of task-relevant inputs and outputs” and “off-line cognition.”

“Cognition is time pressured” Cognition is constrained by the requirements of real-time interaction with the environment, such as the representational bottleneck (e.g., Brooks 1991; Clark 1997; Pfeifer and Scheier 1999).

“We off-load cognitive work onto the environment” Brooks (1990) formulated a similar claim stating that “the world is its own best model.” Another well-known example is Kirsh and Maglio’s (1994) study of the *Tetris* computer game players’ epistemic actions—that is, decision-preparing movements carried out in the world, rather than in the head.

“The environment is part of the cognitive system” The classical example is Hutchins’s (1995) work on distributed cognition, considering, for example, the instruments in a cockpit as parts of the cognitive system.

“Cognition is for action” A claim made, for example, by Franklin (1995), who argued that minds are the control structures of autonomous agents.

“Off-line cognition is body-based” According to M. Wilson (2002, 625), at the time this claim had received the least attention in the cognitive science literature, although “it may in fact be the best documented and most powerful of the six claims.” An early example is the aforementioned work of Lakoff and Johnson (1980, 1999) who argued that abstract concepts are based on metaphors grounded in bodily activity and experience (such as the English expression “to grasp,” which refers to both manual grasping of physical objects and the more abstract grasping of, for example, ideas or concepts). More recently, the underlying mechanisms have been elaborated in terms of *embodied simulation* accounts of conceptualization and cognition (e.g., Gallese and Lakoff 2005; Gallese 2005).

It might be worth noting here that, in one way or another, all of the above six views deal with the sensorimotor interaction between an agent’s body and its environment, but

none addresses the question of whether the body involved necessarily needs to be physical, biological, and so on. Again, from the perspective of psychology or linguistics—with a focus on human cognition and language—this might be understandable, and the physical or biological nature of the body involved might be considered a nonissue. However, Johnson's (2007, x) account of the development of research on the embodiment of language, which also initially focused on the sensorimotor aspects, indicates that the underlying biological mechanisms were initially somewhat overlooked: "In retrospect I now see that the structural aspects of our bodily interactions with our environments upon which I was focusing were themselves dependent on even more submerged dimensions of bodily understanding. It was an important step to probe below concepts, propositions, and sentences into the sensorimotor processes by which we understand our world, but what is now needed is a far deeper exploration into the qualities, feelings, emotions, and bodily processes that make meaning possible."

11.3 Embodiment in Cognitive Science

The introduction referred to M. Wilson's (2002) general statement that "cognitive processes are deeply rooted in the body's interactions with the world" as a useful first approximation of the fundamental claim of embodied cognition research. Different general notions of embodiment and embodied cognition have already been addressed in the previous section. In this section, let us quickly look at the general theoretical conceptions of embodied cognition in cognitive science and in particular philosophy of mind. The following somewhat simplified diagram from Chemero (2009) provides one useful perspective on the current embodied cognition research landscape (figure 11.1). As Chemero pointed out, there are at least two rather different general theoretical frameworks that are both referred to as "embodied cognitive science." One of these, which Chemero referred to as "radical embodied cognitive science," is grounded in the antirepresentationalist and anticomputationalist traditions of eliminativism, American naturalism (such as the work of James and Dewey), and Gibsonian ecological psychology. The more mainstream version of embodied cognitive science, on the other hand, is derived from the traditional representationalist and computationalist theoretical frameworks that have long dominated cognitive science, and therefore is still more or less compatible with these. In a cognitive robotics context, the latter can be exemplified with the popular notion of symbol/representation grounding (Harnad 1990; Ziemke 1999), whereas the former is closer to the antirepresentationalism advocated by embodied AI researchers such as Brooks (1991) and Beer (1995). As Chemero pointed out, although—or maybe just because—the mainstream version of embodied cognitive science can be considered a "watered-down" version of its more radical counterpart, it currently receives significantly more attention in the cognitive sciences.

Chemero's (2009) formulation of radical embodied cognition can be summarized in the following claims:

1. Representational and computational views of embodied cognition are wrong.
2. Embodied cognition should be explained using a particular set of tools T , including dynamical systems theory.
3. The explanatory tools in set T do not posit mental representations.

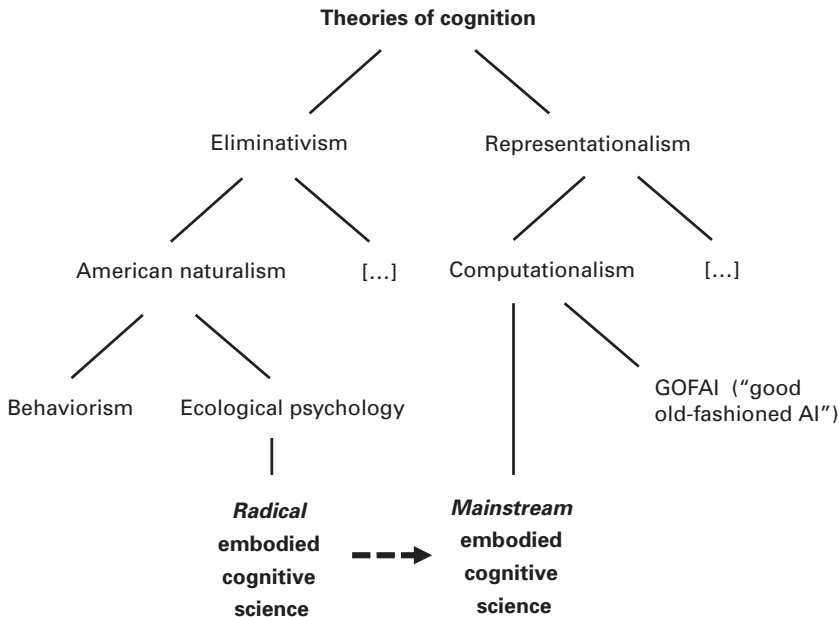


Figure 11.1

Current notions of embodied cognitive science and their historical roots. *Source:* Adapted from Chemero 2009 and Ziemke 2016.

It might be worth noting, though, that Chemero's fundamental and rather strict distinction between representational and antirepresentational approaches to embodied cognition, while offering one useful perspective, should not necessarily be taken to provide some kind of ground truth. First, that seemingly clear-cut distinction obviously hinges on the assumption that there is a more or less wide agreement on what exactly constitutes a representation. This is not necessarily the case (Haselager et al. 2003; Svensson and Ziemke 2005). Hence, it is not difficult to find embodied AI researchers who reject representation in one paper but argue for the grounding of representations in another (e.g., Brooks 1990, 1991). Second, it could be argued that the outright rejection of representation risks throwing out the baby with the bathwater (to use a somewhat dramatic bodily metaphor). That means that while there might very well be good reason to reject the traditional notion of representation, it might be too early, or simply misguided, to reject the notion of representation altogether. Bickhard (1993, 2009), for example, has strongly criticized the traditional notion of representation, which he refers to as *encodingism*, but has developed an interactivist notion of representation, which is much in line with Gibsonian ecological psychology and other elements of radical embodied cognitive science. Similarly, somatic theories of emotion and consciousness, such as the work of Damasio and Panksepp, constitute "a multi-tiered affectively embodied view of mind" (Panksepp 2005, 63), in which representation does play a central role. In this case, however, it is the brain that is considered to "represent" bodily activity, rather than an agent that holds an internal representation of its external environment (cf. Ziemke 2016).

While the controversial issue of representation is certainly too complex to resolve in this chapter, it here suffices to say that, although much early work on embodied cognition

(e.g., Varela et al. 1991) and embodied AI (e.g., Brooks 1991; Beer 1995) was explicitly antirepresentationalist, nowadays much of mainstream embodied cognitive science is more or less representationalist. However, there are at least some accounts of embodied cognition that reject the functionalism/computationalism that characterizes the approaches on the right side of Chemero's diagram without, however, rejecting altogether the notion of representation. Hence, as illustrated in figure 11.2, a revised version of Chemero's diagram, more directly relevant in the context of embodied AI and cognitive robotics, could instead be based on the following distinctions:

- On the one hand, some approaches view cognition as (a) embodied and (b) first and foremost a biological phenomenon; some of these are representationalist in some nontraditional sense (e.g., Damasio, Bickhard), and some of them are antirepresentationalist (where the latter roughly correspond to what Chemero referred to as radical embodied cognitive science).
- On the other hand, there are functionalist approaches—such as traditional cognitive science and GOFAI—that view cognition as first and foremost a computational (and representational) phenomenon; among these we can distinguish between the *computational functionalism* of GOFAI and the *robotic functionalism* (Harnad 1989) that characterizes much of Chemero's notion of mainstream embodied cognitive science (according to which cognition is computational, but its representations need to be grounded in sensorimotor interaction with the environment; cf. Harnad 1990; Pezzulo et al. 2013).

Needless to say, the diagram in figure 11.2 should not necessarily be considered as some kind of ground truth either: First, the picture is not complete (behaviorism, for example, is not included). Second, conceptions of representation, computation, embodiment, and so on

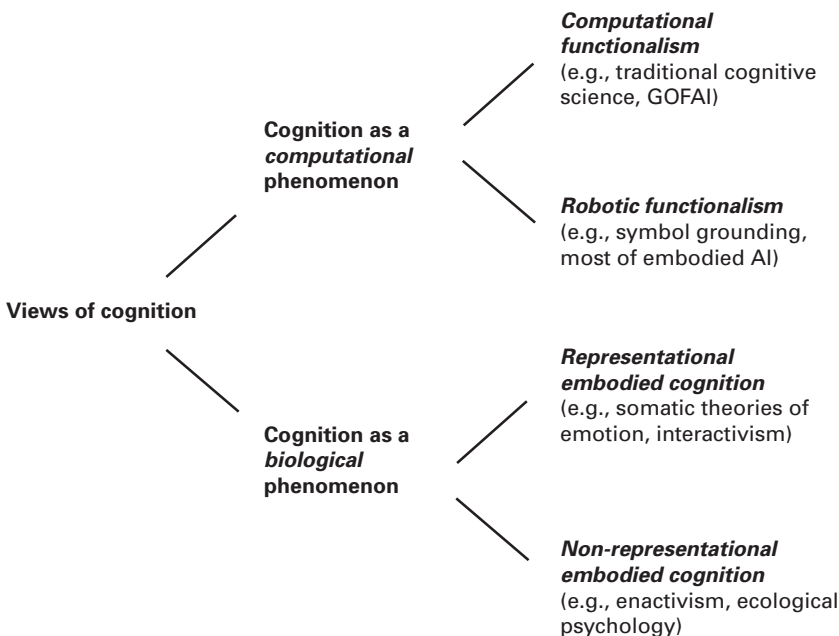


Figure 11.2
Current views of cognition.

obviously vary significantly among researchers. Hence, the reality of the cognitive science and AI research landscape is significantly more complex than either of these diagrams.

11.4 Embodiment in AI

As should be clear from the discussions in previous sections, the issue of embodiment in AI is not straightforward. Many embodied AI researchers, like Brooks (1990) and Steels (1994), emphasize the importance of physical grounding and therefore advocate robotic AI. Others, like Franklin (1997), argue that software agents could be embodied as well if they are situated in an environment (e.g., a search bot searching the internet) and structurally coupled to it. Moreover, many cognitive roboticists in their research practice commonly make use of software simulations of robots and their environments—for example, in order to more quickly train a computational cognitive model in simulation first, which is then later tested on the physical robot. In these cases the computer programs controlling the robots—physical or simulated—are of course for the most part still just as computational as the computer programs of traditional AI. Another case that does not neatly fit into the theoretical categories discussed above are *virtual agents*, as they might appear in video games, for example, or in particular *embodied conversational agents*. Such systems, typically appearing on computer screens, usually have a (simulated) body, used to communicate with their human users, but they typically do not actually use those bodies for sensorimotor interaction with their environment. Hence, while they might *appear* embodied to the people interacting with such systems, in some sense, or to some degree, they really are not embodied in any strong sense.

In fact, most research in embodied AI, although initially often driven by rejections of GOFAI and/or the traditional notion of representation, has been relatively pragmatic in developing the practice of embodied AI, without much concern for philosophical or theoretical distinctions (cf. Ziemke 2004). Based on many years of experience in building embodied AI systems, Pfeifer and colleagues (Pfeifer and Gomez 2005; cf. Pfeifer et al. 2005; Pfeifer and Bongard 2007; Froese and Ziemke 2009) have formulated a number of embodied AI design principles, which together can serve as a characterization of embodied AI as a research field. The first of these are five design procedure principles:

- P1—*synthetic methodology*: aiming for understanding by building.
- P2—*emergence*: systems designed for emergence are often more adaptive.
- P3—*diversity-compliance*: there is a trade-off between exploiting the givens and generating diversity.
- P4—*time perspectives*: three perspectives are required to understand a system's behavior: the “here and now,” its ontogeny (development), and its phylogeny (evolution).
- P5—*frame of reference*: the need to distinguish between observed behavior and underlying mechanisms.

These are complemented by eight agent design principles:

- A1—*three constituents*: an agent, its task, and its ecological niche.
- A2—*complete agents*: focus on embodied, situated, autonomous agents.

- A3—*parallel processes*: asynchronous processes, loosely coupled via the environment.
- A4—*sensorimotor coordination*: self-structured/-generated sensory input.
- A5—*cheap design*: systems exploit their niche and interactions.
- A6—*redundancy*: robustness through overlapping functionalities.
- A7—*ecological balance*: between internal and external complexity.
- A8—*value*: systems have driving forces, development, self-organization.

These five plus eight principles can be seen as guidelines for how to design, build, and/or understand embodied AI systems—where the term “embodied” mainly refers to some form of robotic embodiment and the sensorimotor interaction of internal control and external environment over time. As an elaboration of the *value principle* (A8)—which can be questioned in the case of typical robots that might be argued have no own needs or values, given that they do not have bodies that need to self-maintain, survive, and so on—Froese and Ziemke (2009) have formulated two additional *enactive AI* design principles:

- 1) The system must be capable of generating/maintaining its own systemic identity at some level of description.
- 2) The system must have the capacity to actively regulate its sensorimotor interaction in relation to viability constraints.

The difference or complementarity between embodied and enactive AI principles can be understood in relation to the theoretical distinctions made in previous sections. While the embodied AI principles of Pfeifer and colleagues mainly emphasize physical/sensorimotor embodiment and structural coupling through sensors and actuators, the enactive AI principles additionally emphasize that the organismic embodiment of living systems implies additional constraints and requirements, but also opportunities, that arise from the fact that living bodies need to self-regulate their internal processes and external interactions so as to remain viable, which implies some kind of bodily homeostasis or allostasis (Froese and Ziemke 2009; cf. Damasio and Carvalho 2013; Vernon et al. 2015; Ziemke 2016).

11.5 Embodiment in Cognitive Robotics

Since this chapter is part of a book on cognitive robotics, we will not dwell on trying to define cognitive robotics in detail (see chapter 1 for a more detailed discussion). For the discussion of embodiment, it might be useful, however, to distinguish roughly between the scientific approach and the engineering approach, although in practice they can certainly overlap.

The engineering approach to cognitive robotics could be characterized as the general endeavor to provide robots with cognitive capacities, such as perception, memory, learning, or communication. An example of this is recent work in our lab, and several others, in the European project DREAM (<https://dream2020.github.io/DREAM/>), whose aim it was to develop humanoid robots that could interact with kids with autism as part of psychological therapy, with the goal of teaching social interaction skills, such as joint attention, turn-taking, and imitation (e.g., Cao et al. 2019). This is an example of an engineering

approach because the mechanisms underlying the robots' cognitive and interactive capacities, for the most part, were not *based on* models of human cognition, although they were of course tailor-made to match the cognitive and interactive capacities of the children involved. Hence, the role of the robot's physical embodiment, much like in the case of the embodied conversational agents mentioned above, is not so much that it is fundamental to the robot's cognitive processes as such but rather that the embodiment plays a crucial role in the kids' embodied social interaction with the robot. For example, the robots needed to be able to perceive the same objects, to observe the kids' behavior, to act (e.g., point to objects), and to communicate (e.g., talk) in ways that the kids could understand.

The (cognitive-) scientific approach to cognitive robotics, on the other hand, could be characterized as the use of robotic models for the express purpose of understanding the mechanisms underlying the cognitive and behavioral capacities of humans and/or other animals. Hence, the contribution of cognitive robotics to the study of embodied cognition lies in building robotic models that help elucidate the many ways in which "cognitive processes are deeply rooted in the body's interactions with the world"—to get back to M. Wilson's (2002) characterization of embodied cognition that we used in the introduction. We can further roughly distinguish between minimalist approaches, which usually try to model general mechanisms or principles underlying cognition and behavior, and more complex approaches, which usually try to build specific models, in many cases aiming to replicate data observed in human or animal experiments.

Some of our own work in evolutionary robotics can be used to illustrate the minimalist approach: Ziemke and Thieme (2002), for example, presented experiments using an evolutionary-robotics methodology (Nolfi and Floreano 2000, chap. 4). Simple simulated wheeled robots were evolved to deal with delayed-response tasks that required "memory" of where light sources had previously been perceived in order to find a goal location in a maze. Delayed response tasks are the classical paradigm in psychology for studying working memory in humans and other animals (Malloy 2011). The point of the simple robotic model in this case was not to replicate the body or the data from some specific animal experiment. It was to illustrate how the embodied cognitive capacity (memory) required to solve the delayed response task (i.e., to "remember" where the light was perceived) could emerge from the interplay of the robot's minimal internal mechanisms and its sensorimotor interaction with its environment, rather than from some explicit internal representation in the traditional sense. The "embodiment" of the robotic agents used in such experiments is often intentionally reduced to a bare minimum, which makes it easy to analyze in detail the interaction of internal and external mechanisms (cf. Ziemke 2003a, 2005), which is of course not possible to do in equivalent experiments with animals. Another example of interacting and coadapting agents would be our work (Buason et al. 2005) on evolving robotic agents in a predator-prey scenario. Here predators and prey were given the opportunity to coevolve—that is, adapt to each other, over a series of generations. In a nutshell, the results showed several effects that have also been observed in natural predator-prey coevolution, such as the fact that predators tended to evolve a narrow field of view (suited to pursue the prey in front of them), whereas the prey evolved a significantly wider field of view (suited to detect both obstacles in front of them and predators behind them). Like the above delayed response tasks, this is another example of minimally cognitive behavior (Beer 1996; Barandiaran and Moreno 2006) because

predators and prey had to “remember” each other whenever they temporarily lost track of each other. Again, the “embodiment” in these simulation experiments was intentionally kept minimal—a number of sensors on a simulated, simple circular robot body—and the experiments did not replicate data from some *specific* experiment or species but rather provided insight into the *general* mechanisms of predator-prey coadaptation of body and behavior.

An example of a more complex cognitive robotics experiment aiming to model specific aspects of human cognition, and to also replicate human experimental data, comes from the work of Morse et al. (2015). Using a full-scale humanoid robot, they replicated infant studies investigating the role of bodily posture in how infants learn mappings between words and objects (see chapter 20). The robot model was used to test the hypothesis that a body-centric spatial location, and thus its momentary posture, is used to bind the multimodal features of visual objects and their names. The robot model was shown to replicate data from infant studies and generate novel predictions, which were then tested in new infant studies. This model showed how the memory of name-object mappings, used in new spatial locations, can emerge through the body’s momentary disposition in space.

Hence, although the above robotic models differ radically in the complexity of the robotic embodiment used (simple simulated wheeled robots vs. a physical humanoid), they all address how a cognitive capacity such as memory can emerge—in an embodied manner—from an agent’s sensorimotor interactions with its environment over time.

A negative take on the somewhat perplexing diversity of “embodiments” used by researchers in embodied AI and cognitive robotics would be that 1) researchers simply have not come to any significant agreement on what embodiment “is,” and/or 2) all existing embodied AI systems are still at best very limited versions of the real thing—that is, human or other living bodies. Dreyfus (2007), for example, argued the latter point in his critique of embodied AI, stating that attempts to model human cognition are more or less doomed because they would require “a detailed description of our body and motivations like ours” and that such models “haven’t a chance of being realized in the real world.” However, as we have argued in more detail elsewhere (Froese and Ziemke 2009), the purpose of a model, of course, is usually not to replicate or (re-) instantiate a particular phenomenon in its entirety but rather to help explain it (cf. Morse and Ziemke 2008; Di Paolo and Iizuka 2008). Or, as Froese and Ziemke (2009, 470–471) put it, “It could also be argued that such a detailed modeling approach is not even desirable in the first place since it does not help us to understand why having a particular body allows things in the environment to show up as significant for the agent possessing that body. . . . In other words, instead of blindly modeling the bodies of living beings in as much detail and complexity as possible, it would certainly be preferable to determine the necessary conditions for the constitution of an individual agent with a meaningful perspective on the world.”

As the philosophically minded reader might have noticed by now, this discussion is of course closely related to Searle’s (1980) classical distinction between what he called “weak AI,” the claim that computational models can contribute to our scientific understanding of human cognition, and what he referred to as “strong AI,” the claim that computer models can actually constitute or replicate human or humanlike cognition, consciousness, and so on. Or, in the terms of the discussion in previous sections, if you find yourself sympathizing with functionalist/computationalist approaches to embodied cognition, then you are

likely to think that cognitive robotics could lead to a robotic “strong AI,” at least in theory. The functionalist position has been formulated explicitly by Zlatev (2001, 155), who posited that “a robot with bodily structures, interaction patterns and development similar to those of human beings . . . could possibly recapitulate [human] ontogenesis, leading to the emergence of intentionality, consciousness and meaning.” If, on the other hand, you consider cognition first and foremost a biological phenomenon, then you are likely to think that cognitive robotics is limited to being a form of “weak AI” in Searle’s sense—that is, an approach to the scientific modeling of embodied cognition rather than striving for replication of human or humanlike cognition, in a strong sense. While the term “weak” might sound negative, it should be noted that it is only weak in the sense that it is the weaker—or the more realistic, some would say—of the two claims (weak vs. strong AI). As the examples of cognitive robotics models discussed in this section illustrate, this approach can certainly make strong contributions to our scientific understanding of the mechanisms underlying embodied cognition, especially when used as a complement to other scientific approaches and methodologies.

11.6 Conclusion

We started this chapter by characterizing research on embodiment as guided by the view that “cognitive processes are deeply rooted in the body’s interactions with the world” (M. Wilson 2002, 625). In the context of AI and robotics, it is then natural to ask what kind of body an artifact with certain cognitive capacities might require and how such a system might become grounded (cf. Harnad 1990; Ziemke 1999) in its environment in roughly the way that humans and other animals are. As the discussion in section 11.2 showed, there is not much agreement regarding the first question: What kind of body? The answers range from software agents, to physical robots, to living organisms. In section 11.3 we discussed that we might need to distinguish between at least two fundamentally different conceptions of embodied cognition, which Chemero (2009) referred to as mainstream embodied cognitive science, which has inherited the representationalism and computationalism of traditional cognitive science, and radical embodied cognitive science, which rejects representationalism—at least according to Chemero (2009). We suggested that a more relevant distinction might be between the functionalist view of cognition as first and foremost a computational phenomenon (which at least in principle should be implementable in robots) and the antifunctionalist view of cognition as first and foremost a biological phenomenon (which simply might not be replicable, at least not with current robotic and computational technologies). As discussed in section 11.4, much—if not most—research in embodied AI is somewhat indifferent to such theoretical distinctions (cf. Ziemke 2004), although antirepresentationalism has long been a driving force in early embodied AI. Instead, much embodied AI research has been driven more by the development of a practice of embodied AI—that is, how should we synthesize and analyze embodied forms of (artificial) intelligence—here characterized with a number of embodied and enactive AI design principles. Section 11.5 then discussed the role of embodiment in cognitive robotics and illustrated this with examples of both minimalistic and more complex/human-level robotic models of embodied cognition, in particular how memory can emerge from embodied agent-environment interactions.

From an engineering perspective, the diversity of concepts and approaches discussed here might be somewhat disappointing—we still don't know how to build *Westworld*-type humanlike robots or even *if* we will ever be able to. However, from the scientific perspective of cognitive robotics as an approach to modeling human and animal cognition, the diversity of approaches is actually rather promising because different approaches can complement each other. As the above examples illustrate, cognitive robotics models can function as both

- a complement to theoretical discussions (e.g., helping to clarify overly abstract discussions of “representation” by concrete models of possible underlying mechanisms) and as
- a complement to empirical studies of cognition in humans and animals (e.g., offering better opportunities for the replication and analysis of experiments, as well as a mechanism of hypothesis testing and generation).

A useful step forward for future research on embodied cognition and cognitive robotics, though, might be a clearer distinction between the different “bodies” or perspectives being addressed. From the discussions in this chapter, we can conclude that we have to distinguish between at least

- the *social body*, as it appears to others,
- the *sensorimotor body*, which interacts with the environment,
- the *living body*, which has to self-regulate and self-maintain, and
- the *lived body*, as it is experienced by an agent itself.

In living systems, these four “bodies” are of course really different aspects of the same body, and they roughly correspond to the overlapping perspectives of different disciplines, such as social psychology, behavioral psychology/ethology, biology, and phenomenology, respectively. These multiple bodies also roughly correspond to the first-person (lived body), second-person (social body), and third-person (living, sensorimotor body) perspectives that are fundamental to much of our social cognition and language use.

Moreover, in living systems these multiple bodies are nested in some sense like Russian dolls, to use one final bodily metaphor; that is, the living body motivates and regulates the sensorimotor body's interaction with the environment, which in turn facilitates and manifests the social body and its interactions. In embodied AI and cognitive robotics, however, some of those Russian dolls are usually missing: Most robots have physical/sensorimotor bodies that are not driven by the needs and motivations of an underlying living body. Furthermore, an artificial agent—most obvious in the case of many embodied conversational agents—might appear to have a social body, although it is not necessarily driven and grounded by a sensorimotor body.

While for researchers in cognitive robotics all of this might be relatively transparent, it remains to be seen exactly how this affects the public perception of robotic systems with cognitive and interactive capacities (e.g., Thellman and Ziemke 2020, 2021)—and in particular how it affects people's embodied social interaction with such diversely embodied technologies as humanoid robots, virtual agents, and automated vehicles (cf. Ziemke 2020). Some of those Russian dolls might not be easy to unpack.

Additional Reading and Resources

- *The classical book on the embodied mind and the starting point for enactive cognitive science: Varela, F. J., E. Thompson, and E. Rosch. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.*
- A broad introduction and comprehensive overview of the research area: Shapiro, L. 2010. *Embodied Cognition*. London: Routledge.
- A review paper with a focus on embodied cognition as a biological phenomenon: Ziemke, T. 2016. "The Body of Knowledge: On the Role of the Living Body in Grounding Embodied Cognition." *BioSystems* 148:4–11.

References

- Barandiaran, Xavier, and Alvaro Moreno. 2006. "On What Makes Certain Dynamical Systems Cognitive: A Minimally Cognitive Organization Program." *Adaptive Behavior* 14 (2): 171–185.
- Beer, Randall D. 1995. "A Dynamical Systems Perspective on Agent-Environment Interaction." *Artificial Intelligence* 72 (1–2): 173–215.
- Beer, Randall D. 1996. "Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior." In *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, edited by P. Maes, M. Mataric, J. A. Meyer, J. Pollack, and S. Wilson, 421–429. Cambridge, MA: MIT Press.
- Bickhard, Mark H. 1993. "Representational Content in Humans and Machines." *Journal of Experimental and Theoretical Artificial Intelligence* 5 (4): 285–333.
- Bickhard, Mark H. 2009. "The Biological Foundations of Cognitive Science." *New Ideas in Psychology* 27 (1): 75–84.
- Black, Daniel. 2014. *Embodiment and Mechanisation: Reciprocal Understandings of Body and Machine from the Renaissance to the Present*. Farnham, UK: Ashgate.
- Brooks, R. A. 1990. "Elephants Don't Play Chess." *Robotics and Autonomous Systems*, no. 6: 1–2.
- Brooks, Rodney A. 1991. "Intelligence without Representation." *Artificial Intelligence* 47 (1–3): 139–159.
- Buason, Gunnar, Nicklas Bergfeldt, and Tom Ziemke. 2005. "Brains, Bodies, and Beyond: Competitive Co-evolution of Robot Controllers, Morphologies and Environments." *Genetic Programming and Evolvable Machines* 6 (1): 25–51.
- Cao, Hoang-Long, Pablo G. Esteban, Madeleine Bartlett, Paul Baxter, Tony Belpaeme, Erik Billing, Haibin Cai, et al. 2019. "Robot-Enhanced Therapy: Development and Validation of Supervised Autonomous Robotic System for Autism Spectrum Disorders Therapy." *IEEE Robotics and Automation Magazine* 26 (2): 49–58.
- Chemero, Anthony. 2009. *Radical Embodied Cognitive Science*. MIT Press.
- Clark, Andy. 1997. *Being There*. Cambridge, MA: MIT Press.
- Clark, Andy. 1999. "An Embodied Cognitive Science?" *Trends in Cognitive Science* 3 (9): 345–351.
- Damasio, Antonio R. 1998. "Emotion in the Perspective of an Integrated Nervous System." *Brain Research Reviews* 26 (2–3): 83–86.
- Damasio, Antonio R. 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Boston: Houghton Mifflin Harcourt.
- Damasio, Antonio, and Gil B. Carvalho. 2013. "The Nature of Feelings: Evolutionary and Neurobiological Origins." *Nature Reviews Neuroscience* 14 (2): 143–152.
- Di Paolo, Ezequiel A., and Hiroyuki Iizuka. 2008. "How (Not) to Model Autonomous Behavior." *Biosystems* 91 (2): 409–423.
- Dreyfus, Hubert L. 1979. *What Computers Can't Do*. Cambridge, MA: MIT Press.
- Dreyfus, Hubert L. 2007. "Why Heideggerian AI Failed and How Fixing It Would Require Making It More Heideggerian." *Philosophical Psychology* 20 (2): 247–268.
- Franklin, Stan. 1995. *Artificial Minds*. Cambridge, MA: MIT Press.
- Franklin, Stan. 1997. "Autonomous Agents as Embodied AI." *Cybernetics and Systems* 28 (6): 499–520.

- Froese, Tom, and Tom Ziemke. 2009. "Enactive Artificial Intelligence: Investigating the Systemic Organization of Life and Mind." *Artificial Intelligence* 173 (3–4): 466–500.
- Gallagher, Shaun. 2005. *How the Body Shapes the Mind*. Oxford: Oxford University Press.
- Gallese, Vittorio. 2005. "Embodied Simulation: From Neurons to Phenomenal Experience." *Phenomenology and the Cognitive Sciences* 4 (1): 23–48.
- Gallese, Vittorio, and George Lakoff. 2005. "The Brain's Concepts: The Role of the Sensory-Motor System in Conceptual Knowledge." *Cognitive Neuropsychology* 22 (3–4): 455–479.
- Goldinger, Stephen D., Megan H. Papesh, Anthony S. Barnhart, Whitney A. Hansen, and Michael C. Hout. 2016. "The Poverty of Embodied Cognition." *Psychonomic Bulletin and Review* 23 (4): 959–978.
- Harnad, Stevan. 1989. "Minds, Machines and Searle." *Journal of Experimental and Theoretical Artificial Intelligence* 1 (1): 5–25.
- Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D: Nonlinear Phenomena* 42 (1–3): 335–346.
- Haselager, Pim, André De Groot, and Hans van Rappard. 2003. "Representationalism vs. Anti-representationalism: A Debate for the Sake of Appearance." *Philosophical Psychology* 16 (1): 5–24.
- Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Johnson, Mark. 2007. *The Meaning of the Body: Aesthetics of Human Understanding*. Chicago: University of Chicago Press.
- Kirsh, David, and Paul Maglio. 1994. "On Distinguishing Epistemic from Pragmatic Action." *Cognitive Science* 18 (4): 513–549.
- Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Lakoff, George, and Mark Johnson. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.
- Lindblom, Jessica. 2015. *Embodied Social Cognition*. Vol. 26. Berlin: Springer.
- Lund, Henrik Hautop, Barbara Webb, and John Hallam. 1998. "Physical and Temporal Scaling Considerations in a Robot Model of Cricket Calling Song Preference." *Artificial Life* 4 (1): 95–107.
- Malloy, Paul. 2011. "Delayed Response Tasks." In *Encyclopedia of Clinical Neuropsychology*, edited by J. S. Kreutzer, J. Deluca, and B. Caplan. New York: Springer.
- Maturana, Humberto R., and Francisco J. Varela. 1980. *Autopoiesis and Cognition*. Dordrecht: Reidel.
- Maturana, Humberto R., and Francisco J. Varela. 1987. *The Tree of Knowledge—the Biological Roots of Human Understanding*. Boston: Shambhala.
- Morse, Anthony F., Viridian L. Benitez, Tony Belpaeme, Angelo Cangelosi, and Linda B. Smith. 2015. "Posture Affects How Robots and Infants Map Words to Objects." *PLoS One* 10 (3): e0116012.
- Morse, Anthony F., Carlos Herrera, Robert Clowes, Alberto Montebelli, and Tom Ziemke. 2011. "The Role of Robotic Modelling in Cognitive Science." *New Ideas in Psychology* 29 (3): 312–324.
- Morse, Anthony F., and Tom Ziemke. 2008. "On the Role(s) of Modelling in Cognitive Science." *Pragmatics and Cognition* 16 (1): 37–56.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics*. Cambridge, MA: MIT Press.
- Panksepp, Jaak. 2005. "Affective Consciousness: Core Emotional Feelings in Animals and Humans." *Consciousness and Cognition* 14 (1): 30–80.
- Pezzulo, Giovanni, Lawrence W. Barsalou, Angelo Cangelosi, Martin H. Fischer, Ken McRae, and Michael Spivey. 2013. "Computational Grounded Cognition: A New Alliance between Grounded Cognition and Computational Modeling." *Frontiers in Psychology* 3:612.
- Pfeifer, Rolf, and Josh Bongard. 2007. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, Rolf, and Gabriel Gomez. 2005. "Interacting with the Real World: Design Principles for Intelligent Systems." *Artificial Life and Robotics* 9 (1): 1–6.
- Pfeifer, Rolf, Fumiya Iida, and Josh Bongard. 2005. "New Robotics: Design Principles for Intelligent Systems." *Artificial Life* 11 (1–2): 99–120.
- Pfeifer, Rolf, and Christian Scheier. 1999. *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Quick, Tom, Kerstin Dautenhahn, Chrystopher L. Nehaniv, and Graham Roberts. 1999. "On Bots and Bacteria: Ontology Independent Embodiment." In *Proceedings of the 5th European Conference on Advances in Artificial Life (ECAL '99)*, edited by D. Floreano et al., 339–343. Berlin: Springer-Verlag.
- Riegler, Alexander. 2002. "When Is a Cognitive System Embodied?" *Cognitive Systems Research* 3 (3): 339–348.

- Searle, John. 1980. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3 (3): 417–457.
- Shapiro, Lawrence. 2010. *Embodied Cognition*. London: Routledge.
- Steels, Luc. 1994. "The Artificial Life Roots of Artificial Intelligence." *Artificial Life* 1 (1–2): 75–110.
- Svensson, Henrik, and Tom Ziemke. 2005. "Embodied Representations: What Are the Issues?" In *Proceedings of the Annual Meeting of the Cognitive Science Society* (27).
- Thellman, Sam, and Tom Ziemke. 2020. "Do You See What I See? Tracking the Perceptual Beliefs of Robots." *iScience* 23 (10): 101625.
- Thellman, Sam, and Tom Ziemke. 2021. "The Perceptual Belief Problem: Why Explainability Is a Tough Challenge in Social Robotics." *ACM Transactions on Human-Robot Interaction* 10 (3): 29.
- Thompson, Evan. 2007. *Mind in Life*. Cambridge, MA: Harvard University Press.
- Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
- Vernon, David, Robert Lowe, Serge Thill, and Tom Ziemke. 2015. "Embodied Cognition and Circular Causality: On the Role of Constitutive Autonomy in the Reciprocal Coupling of Perception and Action." *Frontiers in Psychology* 6:1660.
- Wilson, Andrew D., and Sabrina Golonka. 2013. "Embodied Cognition Is Not What You Think It Is." *Frontiers in Psychology* 4:58.
- Wilson, Margaret. 2002. "Six Views of Embodied Cognition." *Psychonomic Bulletin and Review* 9 (4): 625–636.
- Ziemke, Tom. 1999. "Rethinking Grounding." In *Understanding Representation in the Cognitive Sciences*, edited by A. Riegler, M. Peschl, and A. Von Stein. New York: Plenum Press.
- Ziemke, Tom. 2001. "Are Robots Embodied?" In Vol. 85, *Proceedings of the First International Workshop on Epigenetic Robotics: Modelling Cognitive Development in Robotic Systems*, edited by C. Balkenius, J. Zlatev, C. Brezeal, K. Dautenhahn, and H. Kozima, 75–83. Lund, Sweden: University Cognitive Studies.
- Ziemke, Tom. 2003a. "On the Role of Robot Simulations in Embodied Cognitive Science." *AISB Journal* 1 (4): 389–399.
- Ziemke, Tom. 2003b. "What's That Thing Called Embodiment?" In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, edited by R. Alterman and D. Kirsh, 1305–1310. Mahwah, NJ: Lawrence Erlbaum.
- Ziemke, Tom. 2004. "Embodied AI as Science: Models of Embodied Cognition, Embodied Models of Cognition, Or Both?" In *Embodied Artificial Intelligence*, edited by F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi, 27–36. Heidelberg: Springer.
- Ziemke, Tom. 2005. "Cybernetics and Embodied Cognition: On the Construction of Realities in Organisms and Robots." *Kybernetes* 34 (1/2): 118–128.
- Ziemke, Tom. 2016. "The Body of Knowledge: On the Role of the Living Body in Grounding Embodied Cognition." *Biosystems* 148:4–11.
- Ziemke, Tom. 2020. "Understanding Robots." *Science Robotics* 5 (46): eabe2987.
- Ziemke, Tom, and Mikael Thieme. 2002. "Neuromodulation of Reactive Sensorimotor Mappings as a Short-Term Memory Mechanism in Delayed Response Tasks." *Adaptive Behavior* 10 (3–4): 185–199.
- Ziemke, Tom, and Serge Thill. 2014. "Robots Are Not Embodied! Conceptions of Embodiment and Their Implications for Social Human-Robot Interaction." In *Robophilosophy*, edited by Johanna Seibt, Raul Hakli, and Marco Nørskov, 49–53. Amsterdam: IOS Press.
- Ziemke, Tom, Jordan Zlatev, and Roslyn M. Frank. 2006. *Body, Language and Mind. Vol. 1: Embodiment*. Berlin: Mouton De Gruyter.
- Zlatev, Jordan. 2001. "The Epigenesis of Meaning in Human Beings, and Possibly in Robots." *Minds and Machines* 11 (2): 155–195.

12 Ethics of Robotics

Vincent C. Müller

12.1 Introduction

This chapter will provide a comprehensive introduction to the ethics of robotics, with a particular emphasis on the integration of artificial intelligence (AI) and robotics. After the introduction to the field in section 12.1, the main themes are, in section 12.2, ethical issues that arise with robotics systems as *objects* (i.e., tools made and used by humans), where the main sections are privacy, human-robot interaction, employment, and the effects of autonomy, and in section 12.3, robotics systems as *subjects* (i.e., when ethics is for the systems themselves in machine ethics and artificial moral agency). Many of these questions concern the use of AI, so the ethics of AI will play a role in this chapter.

For each section within these themes, we provide a general explanation of the ethical issues, we outline existing positions and arguments, and then we analyze how this plays out with current technologies and finally what policy consequences may be drawn.

12.1.1 Background of the Field

The ethics of robotics is often focused on “concerns” of various sorts—which is a typical response to new technologies. The task of an essay such as this is to analyze the issues and to deflate the nonissues. Some technologies, such as nuclear power, cars, or plastics, have caused ethical and political discussion and significant policy efforts to control the trajectory of these technologies—usually once some damage is done.

The ethics of robotics has seen significant press coverage in recent years, which supports this kind of work but also may end up undermining it: It often talks as though we already knew what would be ethical and as if the issues were just what future technology will bring and what we should do about it. Press coverage thus focuses on considerations of risk, security (Brundage et al. 2018), and the prediction of impact (e.g., on the job market). The result is a discussion of essentially technical problems and on how to achieve the desired outcome. Another result is that much of the current discussion in policy and industry, with its focus on image and public relations—where the label “ethical” is really not much more than the new “green,” is perhaps used for “ethics washing.” For a problem to qualify as a problem for robot ethics would require that we do *not* readily know what is the right thing to do. In this sense, job loss, theft, or killing with a robot are not a problem for ethics, but whether these are permissible under certain circumstances *is* such a problem.

A last caveat is in order for our presentation: The ethics of robotics is a very young field within applied ethics, with significant dynamics but few well-established issues and no authoritative overviews—though surveys for the ethics of robotics exist (Lin, Abney, and Jenkins 2017; Royakkers and van Est 2016; Calo, Froomkin, and Kerr 2016; Tzafestas 2016; European Group on Ethics in Science and New Technologies 2018). So this article cannot just reproduce what the community has achieved thus far but must propose an ordering where little order exists.

12.1.2 A Note on Policy

There is significant public discussion about robot ethics, and there are frequent pronouncements from politicians that the matter requires new policy, but actual technology policy is difficult to plan and to enforce. It can take many forms, from incentives and funding, infrastructure, taxation, or good-will statements to regulation by various actors and the law. Policy for robotics will possibly come into conflict with other aims of technology policy or general policy. One important practical aspect is which agents are involved in the development of a policy and what power structures oversee it.

For people who work in ethics and policy, there is probably a tendency to overestimate the impact and the threats from a new technology and to underestimate how far current regulation can reach (e.g., for product liability). On the other hand, for businesses, the military, and some administrations there is an interest to “talk” and to preserve a good public image but not to “do” anything. Governments, parliaments, associations, and industry circles in industrialized countries have produced reports and white papers in recent years, and some have generated good-will slogans. For a survey, see (Jobin, Ienca, and Vayena 2019).

Though very little actual policy has been produced, there are some notable beginnings. The latest EU policy document suggests “trustworthy AI” should be lawful, ethical, and technically robust and then spells this out as seven requirements: human oversight, technical robustness, privacy and data governance, transparency, fairness, well-being, and accountability (AI HLEG 2019). Much European research now runs under the slogan of “responsible research and innovation” (RRI), and “technology assessment” has been a standard field since the advent of nuclear power. Professional ethics is also a standard field in information technology, and this includes issues that are relevant here. We also expect that much policy will eventually cover specific uses or technologies of robotics, rather than the field as a whole (see Calo 2018; Stahl, Timmermans, and Mittelstadt 2016; Johnson and Verdicchio 2017; Giubilini and Savulescu 2018; Crawford and Calo 2016). The more political angle of technology is often discussed in “science and technology studies” (STS). As books like *The Ethics of Invention* (Jasanoff 2016) show, the concerns are often quite similar to those of ethics (Jacobs et al. 2019).

12.2 Ethics for the Use of Robotics Systems

In this section we outline the ethical issues of the human use of AI and robotics systems that can be more or less autonomous—which means we look at issues that arise with certain uses but would not arise with others. It must be kept in mind, however, that the design of technical

artifacts has ethical relevance for their use (Houkes and Vermaas 2010; Verbeek 2011), so beyond “responsible use,” we also need “responsible design” in this field.

12.2.1 Human-Robot Interaction

Human-robot interaction (HRI) now pays significant attention to ethical matters, to the dynamics of perception from both sides, and to the different interests and the intricacy of the social context, including coworking (e.g., Arnold and Scheutz 2017).

Deception and authenticity

The central questions here often involve whether a robot involves deception, or perhaps violates human dignity or the Kantian requirement of “respect for humanity” (Lin, Abney, and Jenkins 2017). Humans very easily attribute mental properties to objects, and empathize with them, especially when the outer appearance of these objects is similar to that of living beings. This can be used to deceive humans (or animals) into attributing more intellectual or even emotional significance to robots than they deserve. Some parts of humanoid robotics are problematic in this regard (e.g., Hiroshi Ishiguro’s remote-controlled Geminoids), and there are cases that have clearly been deceptive for public relations purposes (e.g., Hanson Robotics’ “Sophia,” with exaggerated statements and even remote control). Of course, some fairly basic constraints of business ethics and law apply to robots too: product safety and liability, or nondeception in advertisement. It appears that these existing constraints take care of many concerns that are raised. There are cases, however, in which HRI has aspects that appear specifically human in ways that can perhaps not be replaced by robots: care, love, and sex.

Example A: Care robots

The use of robots in health care for humans is currently at the level of concept studies in real environments, but it may become a usable technology in a few years and has raised a number of concerns for a dystopian future of dehumanized care (Sharkey and Sharkey 2011; Sparrow 2016). Current systems include robots that support human carers (caregivers)—for example, in lifting patients or transporting material; robots that enable patients to do certain things by themselves, such as eat with a robotic arm; and also robots that are given to patients as company and comfort (e.g., the “Paro” robot seal). For an overview, see (van Wynsberghe 2016; Fosch-Villaronga and Albo-Canals 2019; Nørskov 2017) and for a survey of users Draper et al. (2014).

One reason why the issue of care has come to the fore is that people have argued we will need robots in aging societies. This argument makes problematic assumptions—namely, that with longer life spans people will need more care and that it will not be possible to attract more humans to caring professions. It may also show a bias about age (Jecker 2020). Most importantly, it ignores the nature of automation, which is not simply about replacing humans but about allowing humans to work more effectively. It is not very clear that there really is an issue here since the discussion mostly focuses on the fear of robots dehumanizing care, but the actual and foreseeable robots in care are for the classic automation of technical tasks as assistive robots. They are thus “care robots” only in a behavioral sense of doing what is required, not in the sense that a human “cares” for the patients. It appears that the success of “being cared for” relies on this intentional sense of “care,” which foreseeable robots cannot provide. If anything, the risk of robots in care is the *absence* of

such intentional care—because fewer human carers may be needed. Interestingly, caring for something, even a virtual agent, can be good for the carer themselves (Lee et al. 2019). A system that pretends to care would be deceptive and thus problematic—unless the deception is countered by sufficiently large utility gain (Coeckelbergh 2016). Some robots that pretend to “care” on a basic level are available (Paro seal), and others are in the making. Perhaps feeling cared for by a machine, to some extent, can be progress in some cases?

Example B: Sex robots

Several tech optimists have argued that humans will likely be interested in sex and companionship with robots and feel good about it (Levy 2007). Given the variation of human sexual preferences, including sex toys and sex dolls, this seems very likely: the question is whether such devices should be manufactured and promoted and whether there should be limits to use in this touchy area. It seems to have moved into the mainstream of “robot philosophy” in recent times (Sullins 2012; Danaher and McArthur 2017; Sharkey et al. 2017; Bendel 2018; Devlin 2018).

Humans have long had deep emotional attachments to objects, so perhaps companionship or even love with a predictable android is attractive, especially to people who struggle with actual humans and already prefer dogs, cats, a bird, a computer, or a Tamagotchi. Danaher (2019b) argues against Nyholm and Frank (2017) that this can be true friendship and is thus a valuable goal. It certainly looks like such friendship might increase overall utility, even if lacking in depth. In all these areas, there is an issue of deception since a robot cannot (at present) mean what it says or have feelings for a human. It is well known that humans are prone to attribute feelings and thoughts to entities that behave as if they had sentience and even to clearly inanimate objects that show no behavior at all. Also, paying for deception seems to be an elementary part of the traditional sex industry.

Finally, there are concerns that have often accompanied matters of sex—namely, consent (Frank and Nyholm 2017), aesthetic issues, and worry that humans may be “corrupted” by certain experiences. Old-fashioned though this may seem, human behavior is influenced by experience, and it is likely that pornography or sex robots support the perception of other humans as mere objects of desire, or even as recipients of abuse, and thus ruin a deeper sexual and erotic experience. The Campaign against Sex Robots argues that these devices are a continuation of slavery and prostitution (Richardson 2016).

12.2.2 The Effects of Automation on Employment

It seems clear that AI and robotics will lead to significant gains in productivity and thus overall wealth. The attempt to increase productivity has probably always been a feature of the economy, though the emphasis on “growth” is a modern phenomenon (Harari 2016, 240). However, productivity gains through automation typically mean that fewer humans are required for the same output. This does not necessarily imply a loss of overall employment, however, because available wealth increases and that can increase demand sufficiently to counteract the productivity gain. In the long run, higher productivity in industrial societies has led to more wealth overall. Major labor market disruptions have occurred in the past—for example, farming employed over 60 percent of the workforce in Europe and North America in 1800, while by 2010 it employed about 5 percent in the European Union

and even less in the wealthiest countries (Anonymous 2013). In the twenty years between 1950 and 1970, the number of hired agricultural workers in the UK was reduced by 50 percent (Zayed and Loft 2019). Some of these disruptions lead to more labor-intensive industries moving to places with lower labor cost—this is an ongoing process.

Classic automation replaces human muscle, whereas digital automation replaces human thought or information processing—and unlike physical machines, digital automation is very cheap to duplicate (Bostrom and Yudkovski 2014). It may thus mean a more radical change in the labor market. So the main question is: Is it different, this time? Will the creation of new jobs and wealth keep up with the destruction of jobs? And even if it is *not* different, what are the transition costs, and who bears them? For example, will lower-cost areas suffer and higher-cost areas gain from this development? Do we need to make societal adjustments for a fair distribution of costs and benefits of digital automation?

Responses to the issue of unemployment from robotics and AI have ranged from the alarmed (Frey and Osborne 2013; Westlake 2014) to the neutral (Metcalf, Keller, and Boyd 2016; Calo 2018; Frey 2019) and the optimistic (Brynjolfsson and McAfee 2016; Harari 2016; Danaher 2019a). In principle, the labor market effect of automation seems to be fairly well understood as involving two channels: “(i) the nature of interactions between differently skilled workers and new technologies affecting labor demand and (ii) the equilibrium effects of technological progress through consequent changes in labor supply and product markets” (Goos 2018, 362). What currently seems to happen in the labor market as a result of automation is “job polarization” or the “dumbbell” shape (Goos, Manning, and Salomons 2009): the highly skilled technical jobs are in demand and highly paid, the low-skilled service jobs are in demand and badly paid, but the midqualification jobs in factories and offices—that is, the majority of jobs—are under pressure and reduced because they are relatively predictable and most likely to be automated (Baldwin 2019).

Perhaps enormous productivity gains allow the “age of leisure” to be realized, which Keynes (1930) predicted to occur around 2030, assuming a growth rate of 1 percent per annum? Actually, we have already reached the level he anticipated for 2030, but we are still working—consuming more and inventing ever more levels of organization. Harari explained how this economical development allowed humanity to overcome hunger, disease, and war, and now we aim for immortality and eternal bliss through AI, thus his title *Homo Deus* (Harari 2016, 75).

In general terms, the issue of unemployment is one of how goods in a society should be *justly distributed*. A standard view is that distributive justice should be rationally decided from behind a “veil of ignorance” (Rawls 1971)—that is, as if one does not know what position in a society one would actually be taking (laborer or industrialist, and so on). Rawls thought the chosen principles would then support basic liberties and a distribution that is of greatest benefit to the least-advantaged members of society. It would appear that the robotics economy has three features that make such justice unlikely: First, it operates in a largely unregulated environment where responsibility is often hard to allocate. Second, it operates in markets that have a “winner-takes-all” feature, where monopolies develop quickly. Third, the “new economy” of the digital service industries is based on intangible assets, also called “capitalism without capital” (Haskel and Westlake 2017). This means that it is difficult to control multinational digital corporations that do not rely

on a physical plant in a particular location. These three features seem to suggest that if we leave the distribution of wealth to free market forces, the result would be a heavily unjust distribution. And this is indeed a development that we can already see.

One interesting question that has not received too much attention is whether the development of robotics is environmentally sustainable. Like all computing systems, they produce waste that is very hard to recycle, and they consume vast amounts of energy, especially for the training of machine-learning systems (and even for the mining of cryptocurrency). Again it appears that some agents off-load costs to the general society.

12.2.3 Privacy and Surveillance

There is a general discussion about privacy and surveillance in information technology (e.g., Macnish 2017; Roessler 2017), which mainly concerns the access to private data and data that are personally identifiable. Privacy has several well-recognized aspects—for example, “the right to be left alone,” information privacy, privacy as an aspect of personhood, control over information about oneself, and the right to secrecy (Bennett and Raab 2006). Privacy studies have historically focused on state surveillance by secret services but now include surveillance by other state agents, businesses, and even individuals. The technology has changed massively in the last decades, while regulation has been slow to respond (though there is the GDPR [2016]). The result is an anarchy that is exploited by the most powerful players—sometimes in plain sight, sometimes in hiding.

The digital sphere has widened massively: all data collection and storage are now digital, our lives are more and more digital, most digital data are connected to a single internet, and there is more and more sensor technology around that generates data about nondigital aspects of our lives. At the same time, control over who collects which data, and who has access, is much harder in the digital world than it was in the analog world of paper and telephone calls. Every new technology amplifies the known issues. For example, face recognition in photos and videos allows identification and thus profiling and searching for individuals (Whittaker et al. 2018, 15ff). The result is that “in this vast ocean of data, there is a frighteningly complete picture of us” (Smolan 2016, 1:1), a scandal that still has not received due public attention.

The data trail we leave behind is how our “free” services are paid for, but we are not told about that data collection and its value, and we are manipulated into leaving ever more such data. The primary focus of social media, gaming, and most of the internet in this “surveillance economy” is to gain, maintain, and direct attention—and thus data supply. This surveillance and attention economy is sometimes called “surveillance capitalism” (Zuboff 2019).

Such systems will often reveal facts about us that we ourselves wish to suppress or are not aware of. With the last sentence of his best-selling book *Homo Deus*, Harari (2016) asks about the long-term consequences of AI: “What will happen to society, politics and daily life when non-conscious but highly intelligent algorithms know us better than we know ourselves?”

Robotic devices have not yet played a major role in this area, except for security patrolling, but this will change once they are more common outside of industry environments. Together with the Internet of Things, the “smart” systems (phone, TV, oven, lamp, virtual assistant, home . . .), the “smart city” (Sennett 2018), and “smart governance,” they are set to become part of the data-gathering machinery that offers more detailed data, of different types, in real time, with ever more information.

Privacy-preserving techniques that can conceal the identity of persons or groups to a large extent are now a standard staple in data science; they include (relative) anonymization, access control (plus encryption), and other models in which computation is carried out without access to full unencrypted input data (Stahl and Wright 2018), in the case of “differential privacy” by adding calibrated noise to the output of queries (Dwork et al. 2006; Abowd 2017). While requiring more effort and cost, such techniques can avoid many of the privacy issues. Some companies have also seen better privacy as a competitive advantage that can be leveraged and sold at a price.

12.2.4 Autonomous Systems

Autonomy generally

Several notions of autonomy can be found in the discussion of autonomous systems. A stronger notion is involved in philosophical debates in which autonomy is the basis for responsibility and personhood (Christman 2018). In this context, responsibility implies autonomy, but not inversely, so some systems can have degrees of technical autonomy without raising issues of responsibility. The weaker, more technical, notion of autonomy in robotics is relative and gradual: a system is said to be autonomous with respect to human control to a certain degree (Müller 2012). There is a parallel here to the issues of bias and opacity in AI since autonomy also concerns a power relation: Who is in control, and who is responsible?

Generally speaking, one question is whether autonomous robots raise issues that suggest a revision of present conceptual schemes or whether they just require technical adjustments. In most jurisdictions, there is a sophisticated system of civil and criminal liability to resolve such issues. Technical standards—for example, for the safe use of machinery in medical environments—will likely need to be adjusted. There is already a field of “verifiable AI” for such safety-critical systems and for “security applications.” Bodies like the IEEE and the BSI have produced “standards,” particularly for more technical subproblems, such as data security and transparency. Among the many autonomous systems on land, on water, underwater, in the air, or in space, we discuss two samples: autonomous vehicles and autonomous weapons.

Example A: Autonomous vehicles

Autonomous vehicles hold the promise of reducing the very significant damage that human driving currently causes—with approximately one million humans killed per year, many more injured, the environment polluted, the earth sealed with concrete and tarmac, the cities full of parked cars, and so on. However, there seem to be questions of how autonomous vehicles should behave and how responsibility and risk should be distributed in the complicated system the vehicles operate in. (There is also significant disagreement over how long the development of fully autonomous, or “level 5,” cars [SAE 2015] will actually take.)

There is some discussion of “trolley problems” in this context. In the classic trolley problems (Thompson 1976; Woollard and Howard-Snyder 2016, sect. 2), various dilemmas are presented. The simplest version is that of a trolley train on a track that is heading toward five people and will kill them unless the train is diverted onto a side track. However, on that track is one person who will be killed if the train takes that side track. The example goes back to a remark in (Foot 1967, 6), who discusses a number of dilemma cases in which tolerated and intended consequences of an action differ. Trolley problems are not

supposed to describe actual ethical problems or to be solved with a “right” choice. Rather, they are thought experiments in which choice is artificially constrained to a small, finite number of distinct one-off options and where the agent has perfect knowledge. These problems are used as a theoretical tool to investigate ethical intuitions and theories—especially the difference between actively doing versus allowing something to happen, intended versus tolerated consequences, and consequentialist versus other normative approaches (Kamm and Rakowski 2016). This type of problem has reminded many of the problems encountered in actual driving and in autonomous driving (Lin 2015). It is doubtful, however, that an actual driver or autonomous car will ever have to solve trolley problems (but see Keeling 2019). While autonomous car trolley problems have received a lot of media attention (Awad et al. 2018), they do not seem to offer anything new to either ethical theory or to the programming of autonomous vehicles.

The more common ethical problems in driving, such as speeding, risky overtaking, not keeping a safe distance, and more are classic problems of pursuing personal interest versus the common good. The vast majority of these are covered by legal regulations on driving. Programming the car to drive “by the rules” rather than “by the interest of the passengers” or “to achieve maximum utility” is thus deflated to a standard problem of programming ethical machines (see section 3.1). There are probably additional discretionary rules of politeness and interesting questions on when to break the rules (Lin 2015), but again this seems to be more a case of applying standard considerations (rules vs. utility) to autonomous vehicles.

Notable policy efforts in this field include the report by the German Federal Ministry of Transport and Digital Infrastructure (2017), which stresses that *safety* is the primary objective. Rule 10 states, “In the case of automated and connected driving systems, the accountability that was previously the sole preserve of the individual shifts from the motorist to the manufacturers and operators of the technological systems and to the bodies responsible for taking infrastructure, policy and legal decisions” (see 3.2.1). The resulting German and EU laws on licensing automated driving are much more restrictive than their US counterparts, where “testing on consumers” is a strategy used by some companies—without informed consent of the consumers or the possible victims.

Example B: Autonomous weapons

The notion of automated weapons is fairly old: “For example, instead of fielding simple guided missiles or remotely piloted vehicles, we might launch completely autonomous land, sea, and air vehicles capable of complex, far-ranging reconnaissance and attack missions” (DARPA 1983, 1). This proposal was ridiculed as “fantasy” at the time (Dreyfus, Dreyfus, and Athanasiou 1986, ix), but it is now a reality, at least for more easily identifiable targets (missiles, planes, ships, tanks, and so on) but not for human combatants. The main arguments against (lethal) autonomous weapon systems (AWS or LAWS) are that they support extrajudicial killings, take responsibility away from humans, and make wars or killings more likely—for a detailed list of issues see (Lin, Bekey, and Abney 2008, 73–86).

It appears that lowering the hurdle to use such systems (autonomous vehicles, “fire-and-forget” missiles, or drones loaded with explosives) and reducing the probability of being held accountable would increase the probability of their use. The crucial asymmetry in which one side can kill with impunity and thus has few reasons not to do so already

exists in conventional drone wars with remote-controlled weapons (e.g., the US in Pakistan). It is easy to imagine a small drone that searches, identifies, and kills an individual human—or perhaps a type of human. These are the kinds of cases brought forward by the Campaign to Stop Killer Robots and other activist groups. Some seem to be equivalent to saying that autonomous weapons are indeed weapons, and weapons kill, but we still make them in gigantic numbers. On the matter of accountability, autonomous weapons might make the identification and prosecution of the responsible agents more difficult, but this is not clear given the digital records that one can keep, at least in a conventional war. The difficulty of allocating punishment is sometimes called the “retribution gap” (Danaher 2016).

Another question seems to be whether using autonomous weapons in war would make wars worse or perhaps less bad? If robots reduce war crimes and crimes in war, the answer may well be positive and has been used not only as an argument in favor of these weapons (Arkin 2009; Müller 2016) but also as an argument against (Amoroso and Tamburrini 2018). Arguably, the main threat is not the use of such weapons in conventional warfare but in asymmetric conflicts or by nonstate agents, including criminals.

It has also been said that autonomous weapons cannot conform to International Humanitarian Law, which requires observance of the principles of distinction (between combatants and civilians), proportionality (of force), and military necessity (of force) in military conflict (Sharkey 2019). It is true that the distinction between combatants and noncombatants is difficult to discern, but the distinction between civilian and military ships is easy to see—so all this says is that we should not construct and use such weapons if they do violate humanitarian law. Additional concerns have been raised that being killed by an autonomous weapon threatens human dignity, but even the defenders of a ban on these weapons seem to say that these are not good arguments: “There are other weapons, and other technologies, that also compromise human dignity. Given this, and the ambiguities inherent in the concept, it is wiser to draw on several types of objections in arguments against AWS, and not to rely exclusively on human dignity” (Sharkey 2019).

A lot has been made of keeping humans “in the loop” or “on the loop” of military guidance on weapons—these ways of spelling out “meaningful control” are discussed in Santoni de Sio and van den Hoven (2018). There have been discussions about the difficulties of allocating responsibility for the killings of an autonomous weapon, and a “responsibility gap” has been suggested (esp. Sparrow 2007), meaning that neither the human nor the machine may be responsible. On the other hand, we do not assume that for every event there is someone responsible for that event, and the real issue may well be the distribution of risk (Simpson and Müller 2016). Risk analysis (Hansson 2013) indicates it is crucial to identify who is *exposed* to risk, who is a potential *beneficiary*, and who makes the *decisions* (Hansson 2018, 1822–1824).

12.3 Ethics for Robotics Systems

12.3.1 Machine Ethics

Machine ethics is ethics for machines, for “ethical machines,” and for machines as *subjects* rather than for the human use of machines as *objects*. It is often not very clear whether this is supposed to cover all of robot ethics or to be a part of it (Floridi and Saunders 2004;

Moor 2006; Wallach and Asaro 2017; Anderson and Anderson 2011). Sometimes it looks as though there is the (dubious) inference at play here that if machines act in ethically relevant ways, then we need a machine ethics. Accordingly, some use a broader notion: “Machine ethics is concerned with ensuring that the behavior of machines toward human users, and perhaps other machines as well, is ethically acceptable” (Anderson and Anderson 2007, 15). This might include mere matters of product safety, for example. Some of the discussion in machine ethics makes the very substantial assumption that machines can, in some sense, be ethical agents responsible for their actions, or “autonomous moral agents” (see van Wynsberghe and Robbins 2019). The basic idea of machine ethics is now finding its way into actual robotics, where the assumption that these machines are artificial moral agents in any substantial sense is usually not made (Winfield et al. 2019). It is sometimes observed that a robot that is programmed to follow ethical rules can very easily be modified to follow unethical rules (Vanderelst and Winfield 2018).

The idea that machine ethics might take the form of “laws” has famously been investigated by Isaac Asimov (1942), who proposed “three laws of robotics”: “First Law—A robot may not injure a human being or, through inaction, allow a human being to come to harm. Second Law—A robot must obey the orders given it by human beings except where such orders would conflict with the First Law. Third Law—A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.” Asimov then showed in a number of stories how conflicts between these three laws will make it problematic to use them, despite their hierarchical organization.

It is not clear that there is a consistent notion of “machine ethics” since weaker versions are in danger of reducing “having an ethics” to notions that would not normally be considered sufficient (e.g., without “reflection” or even without “action”); stronger notions that move toward artificial moral agents may describe a—currently—empty set.

12.3.2 Artificial Moral Agents

If one takes machine ethics to concern moral agents, in some substantial sense, then these agents can be called “artificial moral agents” having rights and responsibilities. However, the discussion about artificial entities challenges a number of common notions in ethics, and it can be very useful to understand these in abstraction from the human case (cf. Powers and Ganascia, forthcoming; Misselhorn 2020).

Several authors use “artificial moral agent” in a less demanding sense, borrowing from the software “agent” use in which case matters of responsibility and rights will not arise (Allen, Varner, and Zinser 2000). James Moor (2006) distinguishes four types of machine agents: ethical impact agents (example: robot jockeys), implicit ethical agents (example: safe autopilot), explicit ethical agents (example: using formal methods to estimate utility), and full ethical agents (“Can make explicit ethical judgments and generally is competent to reasonably justify them. An average adult human is a full ethical agent”). Several ways to achieve “explicit” or “full” ethical agents have been proposed, via programming it in (operational morality), via “developing” the ethics itself (functional morality), and finally, full-blown morality with full intelligence and sentience (Allen, Smit, and Wallach 2005; Moor 2006). Programmed agents are sometimes not considered “full” agents because they are “competent without comprehension,” just like the neurons in a brain (Dennett 2017; Hakli and Mäkelä 2019).

In some of these discussions, the notion of “moral patient” plays a role: ethical *agents* have responsibilities, while ethical *patients* have rights, because harm to them matters. It seems clear that some entities are patients without being agents—for example, simple animals that can feel pain but cannot make justified choices. On the other hand, it is normally understood that all agents will also be patients (e.g., in a Kantian framework). Usually, being a person is supposed to be what makes an entity a responsible agent, someone who can have duties and be the object of ethical concerns, and such personhood is typically a deep notion associated with free will (Frankfurt 1971; Strawson 2005) and with having phenomenal consciousness. Torrance (2011) suggests “artificial (or machine) ethics could be defined as designing machines that do things which, when done by humans, are criterial of the possession of ‘ethical status’ in those humans”—which he takes to be “ethical *productivity* and ethical *receptivity*”—his expressions for moral agents and patients.

Responsibility for robots

There is broad consensus that accountability, liability, and the rule of law are basic requirements that must be upheld in the face of new technologies (European Group on Ethics in Science and New Technologies 2018, 18), but the issue is how this can be done and how responsibility can be allocated. If the robots act, will they themselves be responsible, liable, or accountable for their actions? Or should the distribution of risk perhaps take precedence over discussions of responsibility?

Traditional distribution of responsibility already occurs: a car maker is responsible for the technical safety of the car, a driver is responsible for driving, a mechanic is responsible for proper maintenance, the public authorities are responsible for the technical conditions of the roads, and so on. In general “the effects of decisions or actions based on AI are often the result of countless interactions among many actors, including designers, developers, users, software, and hardware. . . . With distributed agency comes distributed responsibility” (Taddeo and Floridi 2018, 751). How this distribution might occur is not a problem that is specific to robotics, but it gains particular urgency in this context (Nyholm 2018a, 2018b).

Rights for robots

Some authors have indicated that whether or not current robots must be allocated rights should be seriously considered (Gunkel 2018a, 2018b; Turner 2019; Danaher 2020). This position seems to rely largely on criticism of the opponents and on the empirical observation that robots and other nonpersons are sometimes treated as having rights. In this vein, a “relational turn” has been proposed: If we relate to robots as though they had rights, then we might be well advised not to search whether they “really” do have such rights, but we should assume that they do (Coeckelbergh 2010, 2012, 2018). This raises the question of how far such antirealism or quasi-realism can go and what it means then to say that “robots have rights” in a human-centered approach (Gerdes 2016). On the other side of the debate, Bryson (2010) has insisted with a useful (but admittedly problematic) slogan that “robots should be slaves”—that is, not enjoy rights, though she considers it a possibility (Gunkel and Bryson 2014).

There is a wholly separate issue of whether robots (or other AI systems) should be given the status of “legal entities” or “legal persons”—in the sense in which natural persons but also states, businesses, or organizations are “entities” and can have legal rights and duties. The European Parliament has considered allocating such status to robots in order to deal

with civil liability (EU Parliament 2016; Bertolini and Aiello 2018) but not criminal liability, which is reserved for natural persons. It would also be possible to assign only a certain subset of rights and duties to robots. It has been said that “such legislative action would be morally unnecessary and legally troublesome” because it would not serve the interest of humans (Bryson, Diamantis, and Grant 2017, 273). In environmental ethics there is a long-standing discussion about the legal rights for natural objects like trees (Stone 1972).

It has also been said that the reasons for developing robots with rights, or artificial moral patients, in the future are ethically doubtful (van Wynsberghe and Robbins 2019). In the community of “artificial consciousness” researchers is significant concern about whether it would be ethical to create such consciousness since this would presumably imply ethical obligations to a sentient being—for example, not to harm it and not to end its existence by switching it off. Some authors have called for a “moratorium on synthetic phenomenology” (Bentley et al. 2018, 28f).

12.4 Conclusion

It is remarkable how imagination or a “vision” of robotics and AI has played a central role since the very beginning of the disciplines in the 1950s. And the evaluation of this vision is subject to dramatic change: In a few decades, we went from the slogans “AI is impossible” (Dreyfus) and “AI is just automation” (Lighthill 1973) to “AI will solve all problems” (Kurzweil 1999) and “AI may kill us all” (Bostrom 2014). This created media attention and public relations efforts, but it also raises the problem of how much of this “philosophy and ethics of AI and robotics” is really an imagined technology. As we said at the outset, AI and robotics have raised fundamental questions about what we should do with these systems, what the systems themselves should do, and what risks they have in the long term. They also challenge the human view of humanity as the intelligent and dominant species on Earth. We have seen the issues that have been raised, and we will have to watch technological and social developments closely to catch the new issues early and to develop a philosophical analysis, as well as to debate the traditional problems of philosophy.

Acknowledgments

This chapter has significant overlap with the article by the same author: Müller, Vincent C. 2020. “Ethics of Artificial Intelligence and Robotics.” In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, 1–70. Palo Alto: CSLI, Stanford University. <https://plato.stanford.edu/entries/ethics-ai/>—I am grateful for the comments of many colleagues on that version.

Parts of the work on this article have been supported by the European Commission under the INBOTS project (H2020 grant no. 780073).

Additional Reading and Resources

- Classic book arguing for the existential risk from AI: Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.

- Short and classic introduction to machine ethics: Moor, James H. 2006. “The Nature, Importance, and Difficulty of Machine Ethics.” *IEEE Intelligent Systems* 21 (4): 18–21.
- Textbook on robot ethics: Royakkers, Lambèr, and Rinie van Est. 2016. *Just Ordinary Robots: Automation from Love to War*. Boca Raton: CRC Press; Taylor and Francis.
- Newsletter on AI ethics in Europe (Charlotte Stix): <https://www.charlottestix.com/europeanaiarchive>.

References

- Abowd, John M. 2017. “How Will Statistical Agencies Operate When All Data Are Private?” *Journal of Privacy and Confidentiality* 7 (3): 1–15.
- AI HLEG. 2019. “High-Level Expert Group on Artificial Intelligence: Ethics Guidelines for Trustworthy AI” European Commission. Last modified March 8, 2021. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.
- Allen, Colin, Iva Smit, and Wendell Wallach. 2005. “Artificial Morality: Top-Down, Bottom-Up, and Hybrid Approaches.” *Ethics and Information Technology* 7 (3): 149–155.
- Allen, Colin, Gary Varner, and Jason Zinser. 2000. “Prolegomena to Any Future Artificial Moral Agent.” *Journal of Experimental and Theoretical Artificial Intelligence* 12 (3): 251–261.
- Amoroso, Daniele, and Guglielmo Tamburrini. 2018. “The Ethical and Legal Case against Autonomy in Weapons Systems.” *Global Jurist* 18 (1).
- Anderson, Michael, and Susan Leigh Anderson. 2007. “Machine Ethics: Creating an Ethical Intelligent Agent.” *AI Magazine* 28 (4): 15–26.
- Anderson, Michael, and Susan Leigh Anderson, eds. 2011. *Machine Ethics*. Cambridge: Cambridge University Press.
- Anonymous. 2013. “How Many People Work in Agriculture in the European Union? An Answer Based on Eurostat Data Sources.” *EU Agricultural Economics Briefs* 8.
- Arkin, Ronald C. 2009. *Governing Lethal Behavior in Autonomous Robots*. Boca Raton: CRC Press.
- Arnold, Thomas, and Matthias Scheutz. 2017. “Beyond Moral Dilemmas: Exploring the Ethical Landscape in Hri.” In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction*, 445–452. New York: IEEE.
- Asimov, Isaac. 1942 [1950]. “Runaround: A Short Story.” *Astounding Science Fiction*. Reprinted in *I, Robot*. New York: Gnome Press, 1950, 40ff.
- Awad, Edmond, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. “The Moral Machine Experiment.” *Nature* 563 (7729): 59–64.
- Baldwin, Richard. 2019. *The Globotics Upheaval: Globalisation, Robotics and the Future of Work*. London: Weidenfeld and Nicolson.
- Bendel, Oliver. 2018. “Sexroboter aus Sicht der Maschinenethik.” In *Handbuch Maschinenethik*, edited by Oliver Bendel, 1–19. Wiesbaden: Springer Fachmedien Wiesbaden.
- Bennett, Colin J., and Charles Raab. 2006. *The Governance of Privacy: Policy Instruments in Global Perspective*. 2nd ed. Cambridge, MA: MIT Press.
- Bentley, Peter J., Miles Brundage, Olle Häggström, and Thomas Metzinger. 2018. “Should We Fear Artificial Intelligence? In-Depth Analysis.” *European Parliamentary Research Service, Scientific Foresight Unit* 614 (547): 1–40.
- Bertolini, Andrea, and Giuseppe Aiello. 2018. “Robot Companions: A Legal and Ethical Analysis.” *Information Society* 34 (3): 130–140.
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Bostrom, Nick, and Eliezer Yudkovski. 2014. “The Ethics of Artificial Intelligence.” In *The Cambridge Handbook of Artificial Intelligence*, edited by Keith Frankish, 316–334. Cambridge: Cambridge University Press.
- Brundage, Miles, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, et al. 2018. “The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation.” ArXiv preprint: 1802.07228.
- Brynjolfsson, Erik, and Andrew McAfee. 2016. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton.

- Bryson, Joanna J. 2010. "Robots Should Be Slaves." In *Close Engagements with Artificial Companions: Key Social, Psychological, Ethical and Design Issues*, edited by Yorick Wilks, 63–74. Amsterdam: John Benjamins.
- Bryson, Joanna J., Mihailis E. Diamantis, and Thomas D. Grant. 2017. "Of, For, and By the People: The Legal Lacuna of Synthetic Persons." *Artificial Intelligence and Law* 25 (3): 273–291.
- Calo, Ryan. 2018. "Artificial Intelligence Policy: A Primer and Roadmap." *University of Bologna Law Review* 3 (2): 180–218.
- Calo, Ryan, Michael A. Froomkin, and Ian Kerr, eds. 2016. *Robot Law*. Cheltenham: Edward Elgar.
- Christman, John. 2018. "Autonomy in Moral and Political Philosophy." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Palo Alto: Stanford University.
- Coeckelbergh, Mark. 2010. "Robot Rights? Towards a Social-Relational Justification of Moral Consideration." *Ethics and Information Technology* 12 (3): 209–221.
- Coeckelbergh, Mark. 2012. *Growing Moral Relations: Critique of Moral Status Ascription*. London: Palgrave.
- Coeckelbergh, Mark. 2016. "Care Robots and the Future of ICT-Mediated Elderly Care: A Response to Doom Scenarios." *AI and Society* 31 (4): 455–462.
- Coeckelbergh, Mark. 2018. "What Do We Mean by a Relational Ethics? Growing a Relational Approach to the Moral Standing of Plants, Robots and Other Non-humans." In *Plant Ethics*, edited by Angela Kallhoff, Marcello Di Paola, and Maria Schörghener, 110–121. London: Routledge.
- Crawford, Kate, and Ryan Calo. 2016. "There Is a Blind Spot in AI Research." *Nature* 538 (7625): 311–313.
- Danaher, John. 2016. "Robots, Law and the Retribution Gap." *Ethics and Information Technology* 18 (4): 299–309.
- Danaher, John. 2019a. *Automation and Utopia: Human Flourishing in a World without Work*. Cambridge, MA: Harvard University Press.
- Danaher, John. 2019b. "The Philosophical Case for Robot Friendship." *Journal of Posthuman Studies* 3 (1): 5–24.
- Danaher, John. 2020. "Welcoming Robots into the Moral Circle: A Defence of Ethical Behaviorism." *Science and Engineering Ethics* 26: 2023–2049.
- Danaher, John, and Neil McArthur, eds. 2017. *Robot Sex: Social and Ethical Implications*. Cambridge, MA: MIT Press.
- DARPA (Defense Advanced Research Projects Agency). 1983. *Strategic Computing: New Generation Computing Technology, a Strategic Plan for Its Development and Application to Critical Problems in Defense*. October 28, 1983. https://www.nitrd.gov/nitrdgroups/images/3/3a/20040929_strategic_computing.pdf.
- Dennett, Daniel C. 2017. *From Bacteria to Bach and Back: The Evolution of Minds*. New York: W. W. Norton.
- Devlin, Kate. 2018. *Turned On: Science, Sex and Robots*. London: Bloomsbury.
- Draپر, Heather, Tom Sorell, Sandra Bedaf, Dag Sverre Syrdal, Carolina Gutierrez-Ruiz, Alexandre Duclos, and Farshid Amirabdollahian. 2014. "Ethical Dimensions of Human-Robot Interactions in the Care of Older People: Insights from 21 Focus Groups Convened in the UK, France and the Netherlands." In *International Conference on Social Robotics*, edited by M. Beetz, B. Johnston, and M. A. Williams. Cham, Switzerland: Springer.
- Dreyfus, Hubert L. 1992. *What Computers Still Can't Do: A Critique of Artificial Reason*. 2nd ed. Cambridge, MA: MIT Press.
- Dreyfus, Hubert L., Stuart E. Dreyfus, and Tom Athanasiou. 1986. *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. *Calibrating Noise to Sensitivity in Private Data Analysis*. Berlin: Springer.
- EU Parliament. 2016. *Draft Report with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))*. January 27, 2017. https://www.europarl.europa.eu/doceo/document/A-8-2017-0005_EN.html.
- European Group on Ethics in Science and New Technologies. 2018. *Statement on Artificial Intelligence, Robotics and "Autonomous" Systems*. Last modified September 3, 2018. http://ec.europa.eu/research/ege/pdf/ege_ai_statement_2018.pdf.
- Floridi, Luciano, and Jeff W. Saunders. 2004. "On the Morality of Artificial Agents." *Minds and Machines* 14:349–379.
- Foot, Philippa. 1967. "The Problem of Abortion and the Doctrine of the Double Effect." *Oxford Review* 5:5–15.
- Fosch-Villaronga, Eduard, and Jordi Albo-Canals. 2019. "'I'll Take Care of You,' Said the Robot: Reflecting upon the Legal and Ethical Aspects of the Use and Development of Social Robots for Therapy." *Paladyn, Journal of Behavioral Robotics* 10 (1): 77–93.

- Frank, Lily, and Sven Nyholm. 2017. "Robot Sex and Consent: Is Consent to Sex between a Robot and a Human Conceivable, Possible, and Desirable?" *Artificial Intelligence and Law* 25 (3): 305–323.
- Frankfurt, Harry. 1971. "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68 (1): 5–20.
- Frey, Carl Benedict. 2019. *The Technology Trap: Capital, Labour, and Power in the Age of Automation*. Princeton, NJ: Princeton University Press.
- Frey, Carl Benedict, and Michael A. Osborne. 2013. "The Future of Employment: How Susceptible are Jobs to Computerisation?" Oxford Martin School Working Papers. September 1, 2013. <https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment/>.
- GDPR. 2016. "General Data Protection Regulation: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of such Data, and Repealing Directive 95/46/EC." *Official Journal of the European Union* 119:1–88.
- Gerdes, Anne. 2016. "The Issue of Moral Consideration in Robot Ethics." *SIGCAS Computers and Society* 45 (3): 274–279.
- German Federal Ministry of Transport and Digital Infrastructure. 2017. *Report of the Ethics Commission: Automated and Connected Driving*. Federal Ministry of Transport and Digital Infrastructure. June 2017. https://www.bmvi.de/SharedDocs/EN/publications/report-ethics-commission.pdf?__blob=publicationFile.
- Giubilini, Alberto, and Julian Savulescu. 2018. "The Artificial Moral Advisor: The 'Ideal Observer' Meets Artificial Intelligence." *Philosophy and Technology* 31 (2): 169–188.
- Goos, Maarten. 2018. "The Impact of Technological Progress on Labour Markets: Policy Challenges." *Oxford Review of Economic Policy* 34 (3): 362–375.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2009. "Job Polarization in Europe." *American Economic Review* 99 (2): 58–63.
- Gunkel, David J. 2018a. "The Other Question: Can and Should Robots Have Rights?" *Ethics and Information Technology* 20 (2): 87–99.
- Gunkel, David J. 2018b. *Robot Rights*. Cambridge, MA: MIT Press.
- Gunkel, David J., and Joanna Bryson. 2014. "Introduction to the Special Issue on Machine Morality: The Machine as Moral Agent and Patient." *Philosophy and Technology* 27(1): 5–8.
- Hakli, Raul, and Pekka Mäkelä. 2019. "Moral Responsibility of Robots and Hybrid Agents." *Monist* 102 (2): 259–275.
- Hansson, Sven Ove. 2013. *The Ethics of Risk: Ethical Analysis in an Uncertain World*. New York: Palgrave Macmillan.
- Hansson, Sven Ove. 2018. "How to Perform an Ethical Risk Analysis (Era)." *Risk Analysis* 38 (9): 1820–1829.
- Harari, Yuval Noah. 2016. *Homo Deus: A Brief History of Tomorrow*. New York: Harper.
- Haskel, Jonathan, and Stian Westlake. 2017. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton, NJ: Princeton University Press.
- Houkes, Wybo, and Pieter E. Vermaas. 2010. *Technical Functions: On the Use and Design of Artefacts*. Berlin: Springer.
- Jacobs, An, Lynn Tytgat, Michel Maus, Romain Meeusen, and Bram Vanderborght, eds. 2019. *Homo Roboticus: 30 Questions and Answers on Man, Technology, Science and Art*. Brussels: ASP.
- Jasanoff, Sheila. 2016. *The Ethics of Invention: Technology and the Human Future*. New York: W. W. Norton.
- Jecker, Nancy S. 2020. *Ending Midlife Bias: New Values for Old Age*. New York: Oxford University Press.
- Jobin, Anna, Marcello Ienca, and Effy Vayena. 2019. "The Global Landscape of AI Ethics Guidelines." *Nature Machine Intelligence* 1 (9): 389–399.
- Johnson, Deborah G., and Mario Verdicchio. 2017. "Reframing AI Discourse." *Minds and Machines* 27 (4): 575–590.
- Kamm, Frances Myrna, and Eric Rakowski, eds. 2016. *The Trolley Problem Mysteries*. New York: Oxford University Press.
- Keeling, Geoff. 2019. "Why Trolley Problems Matter for the Ethics of Automated Vehicles." *Science and Engineering Ethics* 26 (1): 293–307.
- Keynes, John Maynard. 1932. "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*, 358–373. New York: Harcourt Brace.
- Kurzweil, Ray. 1999. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. London: Penguin.

- Lee, Minha, Sander Ackermans, Nena van As, Hanwen Chang, Enzo Lucas, and Wijnand Ijsselstein. 2019. "Caring for Vincent: A Chatbot for Self-Compassion." *CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, no. 702. <https://doi.org/10.1145/3290605.3300932>.
- Levy, David. 2007. *Love and Sex with Robots: The Evolution of Human-Robot Relationships*. New York: Harper.
- Lighthill, James. 1973. "Artificial Intelligence: A General Survey." In *Artificial Intelligence: A Paper Symposium*, 1–21. London: Science Research Council.
- Lin, Patrick. 2015. "Why Ethics Matters for Autonomous Cars." In *Autonomous Driving*, edited by M. Maurer et al., 69–85. Berlin: Springer.
- Lin, Patrick, Keith Abney, and Ryan Jenkins, eds. 2017. *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*. New York: Oxford University Press.
- Lin, Patrick, George Bekey, and Keith Abney. 2008. "Autonomous Military Robotics: Risk, Ethics, and Design." US Department of Navy, Office of Naval Research. http://ethics.calpoly.edu/onr_report.pdf.
- Macnish, Kevin. 2017. *The Ethics of Surveillance: An Introduction*. London: Routledge.
- Metcalfe, Jacob, Emily F. Keller, and Danah Boyd. 2016. "Perspectives on Big Data, Ethics, and Society." Council for Big Data, Ethics, and Society. May 23, 2016. <https://bdes.datasociety.net/council-output/perspectives-on-big-data-ethics-and-society/>.
- Misselhorn, Catrin. 2020. "Artificial Systems with Moral Capacities? A Research Design and Its Implementation in a Geriatric Care System." *Artificial Intelligence* 278:103179.
- Moor, James H. 2006. "The Nature, Importance, and Difficulty of Machine Ethics." *IEEE Intelligent Systems* 21 (4): 18–21.
- Müller, Vincent C. 2012. "Autonomous Cognitive Systems in Real-World Environments: Less Control, More Flexibility and Better Interaction." *Cognitive Computation* 4 (3): 212–215.
- Müller, Vincent C. 2016. "Autonomous Killer Robots Are Probably Good News." In *Drones and Responsibility: Legal, Philosophical and Socio-technical Perspectives on the Use of Remotely Controlled Weapons*, edited by Ezio Di Nucci and Filippo Santoni De Sio, 67–81. London: Ashgate.
- Nørskov, Marco, ed. 2017. *Social Robots*. London: Routledge.
- Nyholm, Sven. 2018a. "Attributing Agency to Automated Systems: Reflections on Human–Robot Collaborations and Responsibility-Loci." *Science and Engineering Ethics* 24 (4): 1201–1219.
- Nyholm, Sven. 2018b. "The Ethics of Crashes with Self-Driving Cars: A Roadmap, II." *Philosophy Compass* 13 (7): e12506.
- Nyholm, Sven, and Lily Frank. 2017. "From Sex Robots to Love Robots: Is Mutual Love with a Robot Possible?" In *Robot Sex: Social and Ethical Implications*, edited by John Danaher and Neil McArthur, 219–243. Cambridge, MA: MIT Press.
- Powers, Thomas M., and Jean-Gabriel Ganascia. Forthcoming. "The Ethics of the Ethics of AI." In *Oxford Handbook of Ethics of Artificial Intelligence*, edited by Markus D. Dubber, Frank Pasquale, and Sunnit Das.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Belknap Press.
- Richardson, Kathleen. 2016. "Sex Robot Matters: Slavery, the Prostituted, and the Rights of Machines." *IEEE Technology and Society* 35 (2).
- Roessler, Beate. 2017. "Privacy as a Human Right." *Proceedings of the Aristotelian Society* 2 (117).
- Royakkers, Lambèr, and Rinie van Est. 2016. *Just Ordinary Robots: Automation from Love to War*. Boca Raton: CRC Press, Taylor and Francis.
- SAE International. 2015. "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles." *SAE Recommended Practice J3016_201806*.
- Santoni De Sio, Filippo, and Jeroen van den Hoven. 2018. "Meaningful Human Control over Autonomous Systems: A Philosophical Account." *Frontiers in Robotics and AI* 5 (15).
- Sennett, Richard. 2018. *Building and Dwelling: Ethics for the City*. London: Allen Lane.
- Sharkey, Amanda. 2019. "Autonomous Weapons Systems, Killer Robots and Human Dignity." *Ethics and Information Technology* 21 (2): 75–87.
- Sharkey, Amanda, and Noel Sharkey. 2011. "The Rights and Wrongs of Robot Care." In *Robot Ethics: The Ethical and Social Implications of Robotics*, edited by Patrick Lin, Keith Abney, and George Bekey, 267–282. Cambridge, MA: MIT Press.
- Sharkey, Noel, Aimee Van Wynsberghe, Scott Robbins, and Eleanor Hancock. 2017. "Report: Our Sexual Future with Robots." Responsible Robotics. July 5, 2017. <https://responsiblerobotics.org/2017/07/05/fr-report-our-sexual-future-with-robots/>.

- Simpson, Thomas W., and Vincent C. Müller. 2016. "Just War and Robots Killings." *Philosophical Quarterly* 66 (263): 302–322.
- Smolan, Sandy. 2016. "The Human Face of Big Data." PBS documentary. 56 mins.
- Sparrow, Rob. 2007. "Killer Robots." *Journal of Applied Philosophy* 24 (1): 62–77.
- Sparrow, Rob. 2016. "Robots in Aged Care: A Dystopian Future." *AI and Society* 31 (4): 1–10.
- Stahl, Bernd Carsten, Job Timmermans, and Brent Daniel Mittelstadt. 2016. "The Ethics of Computing: A Survey of the Computing-Oriented Literature." *ACM Computing Surveys* 48/4 (55): 1–38.
- Stahl, Bernd Carsten, and David Wright. 2018. "Ethics and Privacy in AI and Big Data: Implementing Responsible Research and Innovation." *IEEE Security and Privacy* 16 (3).
- Stone, Christopher D. 1972. "Should Trees Have Standing—toward Legal Rights for Natural Objects." *Southern California Law Review* 2:450–501.
- Strawson, Galen. 2005. *Free Will*. London: Routledge. Last modified February 29, 2004. <http://www.rep.routledge.com/article/v014>.
- Sullins, John P. 2012. "Robots, Love, and Sex: The Ethics of Building a Love Machine." *IEEE Transactions on Affective Computing* 3 (4): 398–409.
- Taddeo, Mariarosaria, and Luciano Floridi. 2018. "How AI Can Be a Force for Good." *Science* 361 (6404): 751–752.
- Thompson, Judith Jarvis. 1976. "Killing, Letting Die and the Trolley Problem." *Monist* 59:204–217.
- Torrance, Steve. 2011. "Machine Ethics and the Idea of a More-than-Human Moral World." In *Machine Ethics*, edited by Michael Anderson and Susan Leigh Anderson, 115–137. Cambridge: Cambridge University Press.
- Turner, Jacob. 2019. *Robot Rules: Regulating Artificial Intelligence*. Berlin: Springer.
- Tzafestas, Spyros G. 2016. *Roboethics: A Navigating Overview*. Berlin: Springer.
- Vanderelst, Dieter, and Alan Winfield. 2018. "The Dark Side of Ethical Robots." In *AIES '18: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 317–322. <https://doi.org/10.1145/3278721.3278726>.
- van Wynsberghe, Aimee. 2016. *Healthcare Robots: Ethics, Design and Implementation*. London: Routledge.
- van Wynsberghe, Aimee, and Scott Robbins. 2019. "Critiquing the Reasons for Making Artificial Moral Agents." *Science and Engineering Ethics* 25 (3): 719–735.
- Verbeek, Peter-Paul. 2011. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press.
- Wallach, Wendell, and Peter M. Asaro, eds. 2017. *Machine Ethics and Robot Ethics*. London: Routledge.
- Westlake, Stian, ed. 2014. *Our Work Here Is Done: Visions of a Robot Economy*. London: Nesta.
- Whittaker, Meredith, Kate Crawford, Roel Dobbe, Genevieve Fried, Elizabeth Kazianas, Varoon Mathur, Sarah Myers West, Rashida Richardson, and Jason Schultz. 2018. "AI Now Report 2018." New York University. https://ainowinstitute.org/ai_now_2018_report.html.
- Winfield, Alan F., Katina Michael, Jeremy Pitt, and Vanessa Evers. 2019. "Machine Ethics: The Design and Governance of Ethical AI and Autonomous Systems." *Proceedings of the IEEE* 107 (3): 509–517.
- Woollard, Fiona, and Frances Howard-Snyder. 2016. "Doing vs. Allowing Harm." In *Stanford Encyclopedia of Philosophy* Fall 2021 edition, edited by Edward N. Zalta. Palo Alto: Stanford University. <https://plato.stanford.edu/archives/fall2021/entries/doing-allowing/>.
- Zayed, Yago, and Philip Loft. 2019. "Agriculture: Historical Statistics." *House of Commons Briefing Paper* 3339:1–19.
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.

III BEHAVIORAL AND COGNITIVE CAPABILITIES

13 Intrinsic Motivations for Open-Ended Learning

Gianluca Baldassarre

13.1 Introduction

Cognitive robotics and machine learning are producing a growing amount of works on *intrinsic motivations* (IMs) and *open-ended learning*. IMs, often contrasted to *extrinsic motivations* (EMs) that in animals are directed to satisfy biological needs such as hunger and thirst, refer to processes such as curiosity, surprise, novelty, and success at accomplishing one's own goals (Barto et al. 2004; Oudeyer et al. 2007; Baldassarre 2011; Baldassarre and Mirolli 2013). Open-ended learning refers to robots and agents that, similarly to the early development of humans (Weng et al. 2001; Lungarella et al. 2003), undergo prolonged periods of learning wherein they autonomously acquire knowledge and skills that might be useful to later solve tasks given by the user (Seepanomwan et al. 2017; Doncieux et al. 2018).

IMs are very important for robotics and machine learning because they can drive the autonomous open-ended learning of robots and machines by requiring little or no human intervention to furnish guidance in terms of data sets, behaviors to imitate, tasks, reward functions, and goals. Moreover, they allow the construction of robots and machines able to robustly operate in cluttered and unstructured environments posing challenges that cannot be anticipated at design time and preventing the possibility of programming behaviors in advance. Consider, for example, *service robots* that have to operate in warehouses, offices, houses, and health-care environments and in the fields of construction, agri-food, and space. Despite this importance, IMs are a subtle concept, as they come in different types, involve both functions (“what they are for”) and mechanisms (“how do they work”), and can be mixed in various ways in the components of cognitive systems and robot controllers. This tends to generate quite a lot of confusion and to make it difficult to choose between the different available solutions when implementing robots and machines. This chapter addresses this problem in two ways. First (section 13.2), it provides computationally driven conceptual grids to define IMs by contrasting them with EMs and then to classify different types of IMs based on their possible functions and mechanisms, in particular by referring to three main classes of IMs here referred to as *epistemic intrinsic motivations* (eIMs). Second (section 13.3), it presents a selection of example models from cognitive robotics and machine learning to show how different IMs can be used to face

different computational problems. The work concludes (section 13.4) by presenting some of the open challenges of the research on IMs.

13.2 Conceptual Grids: Mechanisms and Functions of Extrinsic and Intrinsic Motivations and Classes of (Epistemic) Intrinsic Motivations

The concept of IM has been proposed and developed within the psychological literature to overcome the difficulties of the behaviorist theory on learning and drives (e.g., Skinner 1938; Hull 1943), in particular to explain why animals spontaneously engage in puzzles (Harlow 1950) or can be instrumentally conditioned to produce particular responses on the basis of apparently neutral stimuli (e.g., a sudden light onset; Kish 1955), as happens with “standard” primary rewards (e.g., food). Subsequent proposals highlighted how the properties of certain stimuli can trigger animals’ exploration and guide their learning processes—for example, when the stimuli are complex, unexpected, or in general surprising (Berlyne 1966). Another important thread of psychological research highlighted the importance that action plays in IMs—for example, in relation to the motivation coming from the fact that an agent manages to affect the environment with its behavior (*effectance*; White 1959) or can autonomously set its own goals and master their achievement (Ryan and Deci 2000). IMs involving actions are also related to sensorimotor contingencies studied by psychology and involving the mechanisms underlying the keen interest of animals and humans for the effects of their own actions (Polizzi di Sorrentino et al. 2014; Taffoni et al. 2014; Jacquey et al. 2019).

Within the computational sciences, Schmidhuber (1991a, 1991b) was the first to present a computational operationalization of some IM mechanisms (in particular *prediction-based IMs*; see below), and Barto et al. (2004) settled the fundamental link between IMs (in particular *competence-based IMs*; see below) and reinforcement-learning (RL) methods (Sutton and Barto 2018). These initial ideas were first developed within the developmental robotics scientific community (with works in the *IEEE Transactions on Autonomous Mental Development* journal, the International Conference on Development and Learning, and the Epigenetic Robotics Conference; Zlatev and Balkenius 2001; Lungarella et al. 2003; Oudeyer et al. 2007; Schembri et al. 2007; Doya and Taniguchi 2019), and more recently have been developed within the autonomous/cognitive robotics and machine-learning community (e.g., Bellemare 2016; Nair et al. 2018), in particular driven by the success of deep neural networks and RL (Goodfellow et al. 2017; Sutton and Barto 2018).

We now focus on understanding and defining these concepts more in detail and furnish conceptual grids on them. These grids are grounded in two perspectives from which one can look at cognitive processes (Tinbergen 1963; Marr and Poggio 1976): 1) the *computational functions* they serve—that is, the problems they solve: these indicate the possible “uses” for which they might be employed within an overall cognitive/robotic system; 2) the *mechanisms*, or algorithms: these refer to the information operations used to accomplish the functions. Some specifications are due on how the terms “functions” and “mechanisms” are used here. First, for animals “function” refers to *adaptive function*—that is, the utility of certain elements of intelligence, such as an IM, for the animal’s biological fitness. For robots, “function” refers to the utility of a certain element of the robot’s intelligence for the robot’s *user*. Second, as with the functions in a computer program, “functions” can

be organized at multiple hierarchical levels: from the highest level just mentioned (“biological fitness”; “utility for the user”) to lower levels. For example, “moving an object as desired” can be further decomposed into lower-level functions such as “recognizing the object position” and “issuing suitable motor commands.” Thus, a function can be seen as realized through a mechanism, but this mechanism in turn can be seen as a function to be realized with lower-level mechanisms. This downward decomposition can continue until some mechanisms are reached that are (arbitrarily) considered *primitive* for a given analysis.

13.2.1 Extrinsic and Intrinsic Motivations

What are motivations? Motivations are an element of intelligence having at least three important functions (for organisms; cf. Panksepp 1998): 1) *selection* drives the system to select a behavior, among alternative available ones, to attend the most important current needs/goals; 2) *energy* establishes the amount of energy invested in executing the selected behavior; 3) *learning* generates learning signals to change behavior. This chapter considers in particular the first and third functions of motivations. For example, we will see how IMs can drive an agent to move to some areas of the environment in navigation tasks (behavior selection) or can produce the reward signals for RL processes (production of learning signals).

What are *intrinsic motivations*? When initially studied in psychology, IMs were defined as motivations driving the performance of behavior “for its own sake”—that is, without any direct apparent purpose (Berlyne 1966). Although useful to guide intuition, this definition clarifies neither the functions nor the mechanisms of IMs. A more operational definition proposed here is that *intrinsic motivations are processes that can drive the acquisition of knowledge and skills in the absence of extrinsic motivations* (cf. Baldassarre 2011). IMs are hence best understood by contrasting them to *extrinsic motivations* (EMs). Table 13.1 highlights the main differences between EMs and a very important subset of IMs we will call *epistemic intrinsic motivations* (eIMs). In Baldassarre (2011) eIMs were considered to be IMs tout court, but here we recognize that they do not cover the full spectrum of

Table 13.1
Main features of extrinsic and (epistemic) intrinsic motivations (eIMs)

	Extrinsic motivations (EMs)	(Epistemic) intrinsic motivations (eIMs)
Function	Organisms: acquisition of <i>material resources</i> . Robots: accomplishment of <i>user’s goals</i> .	Acquisition of <i>knowledge and skills</i> .
Mechanism	Organisms: measure the acquisition of <i>material resources</i> by getting information on their levels/changes <i>from body and resource monitoring</i> . Robots: measure the level/change of accomplishment of the <i>user’s goals</i> .	Measure the acquisition of knowledge and skills by getting <i>information</i> on their levels/changes in <i>other parts of the brain (organisms) or controller (robots)</i> .
Time of contribution to the “ultimate” (extrinsic) function	<i>Immediately</i> : when the material resource is acquired and used (organisms); when the user’s goals are accomplished (robots).	<i>Later</i> : when the acquired knowledge and skills are used to acquire resources (organisms) or to accomplish the user’s goals (robots).
“Time signature” of the motivation	They tend to <i>go away</i> when the related resources are acquired and to <i>come back</i> when there is a lack of those resources.	They tend to <i>go away for good</i> when the related pieces of knowledge/skills are acquired.

IMs because, as we shall see, there are some IMs, which we call *other IMs* (oIMs), that are not eIMs. In the table, EMs are contrasted to eIMs because these form the core of IMs and because for their distinctive features they can help to clarify the overall nature of all IMs. The table entries illustrate this in more detail.

Regarding functions, EMs have the overall function of driving behavior and learning to the acquisition of material resources (Baldassarre 2011). For example, the EM of “hunger” drives behavior to look for and ingest food, and when this happens the behavior leading to it is strengthened. Instead, IMs have the overall function of driving behavior and learning toward the acquisition of knowledge and skills (note that “knowledge” also encompasses skills, but here “skills” are referred to explicitly to emphasize the aspects of knowledge more directly linked to action). For example, an IM related to novelty seeking could drive an agent to explore a novel object to learn its appearance, weight, shape, and so on. This function is shared by all IMs, not only by eIMs, as all IMs support the acquisition of knowledge and skills: in other words, all IMs have an epistemic function. In this respect, the term “epistemic motivations” might have been used in place of the term “intrinsic motivations,” which is somehow a misnomer as “intrinsic” suggests “internal” or at best, stretching it, “not directed to external material resources.” However, the term “intrinsic motivations” is kept here for its tradition. Moreover, the term IMs is handy to refer also to oIMs that, contrary to eIMs, are not based on an epistemic mechanism. In this respect, eIMs are the most prototypical IMs as they encompass *both* an epistemic function and an epistemic mechanism, and thus having a term that refers only to them is useful.

In terms of mechanisms, in animals EMs are based on measures of the acquisition of material resources by getting information on their levels/changes in the body or in the environment. For example, hunger, a drive guiding the selection of behaviors related to food seeking, might be triggered when the blood glucose level is low, and a reward-learning signal might be produced when food is ingested. Alternatively, an EM might be related to detecting the presence/availability of resources externally to the body—for example, the presence of a mating companion or the smell of prey in the environment (Baldassarre 2011). In robots, EMs are based on the measure of the accomplishment of the user’s goals; for example, a robot might self-charge its battery to remain operational and bring some objects to the user. Here the terms “extrinsic tasks/goals” will thus be referring to tasks/goals involving the acquisition of material resources or the accomplishment of the user’s goals. Incidentally, notice how EMs are a direct derivation of an evolutionary process not only for animals but also for robots: in animals, the acquisition of material resources is a means to increase biological fitness (number of fertile offspring) and, more specifically, the means for it—that is, survival and reproduction. Similarly, in robots the successful accomplishment of the user’s goals produces a higher chance that the specific features of the robot controller and physical structure are “reproduced,” as they are or in variants, in future robots.

Differently from EMs, eIMs rely on mechanisms that measure knowledge and skills by getting information on their levels/changes in other parts of the brain (for organisms) or in the controller (for robots). Importantly, this implies that an eIM involves the presence of at least three structures and functions inside the brain/controller (figure 13.1): (a) a *source component* that acquires knowledge; (b) an “*IM mechanism*” that measures the level or change of the knowledge of the source component; (c) a “*target component*” that

receives the output of the IM mechanism and uses it to select behavior/energize behavior/drive learning processes. The core of this whole process is (b), the IM mechanism that measures the level or change of knowledge of the source component.

The specification above is very important, as, conceptually, eIMs involve the learning processes and knowledge of two different cognitive/computational components that might be very different in terms of the mechanisms and functions they play within the overall system, and this might make it difficult to recognize them in organisms or to recognize/implement them in robots. In some cases (figure 13.1a), the source component and the target component are the same data structure, in the sense that the IM mechanism detects the knowledge level/change in a component with the function of affecting the learning of the same component (possibly with the mediation of other components; figure 13.1b). For example, the selection of the skill to be trained among many skills to be learned might be based on the competence improvement of the skill itself (e.g., a robot might focus on learning to move one object, rather than on grasping it, if learning the first skill proceeds faster than for the second skill). In other cases (figure 13.1c), the source component and target components are distinct. For example, a component of a robot might detect the novelty of some objects, and this might drive a motor component to explore them with the function of improving its motor ability to manipulate them.

IMs that are not eIMs differ from the latter, as they do not use a learning source component as the origin of the motivation but rather other mechanisms: as anticipated, these

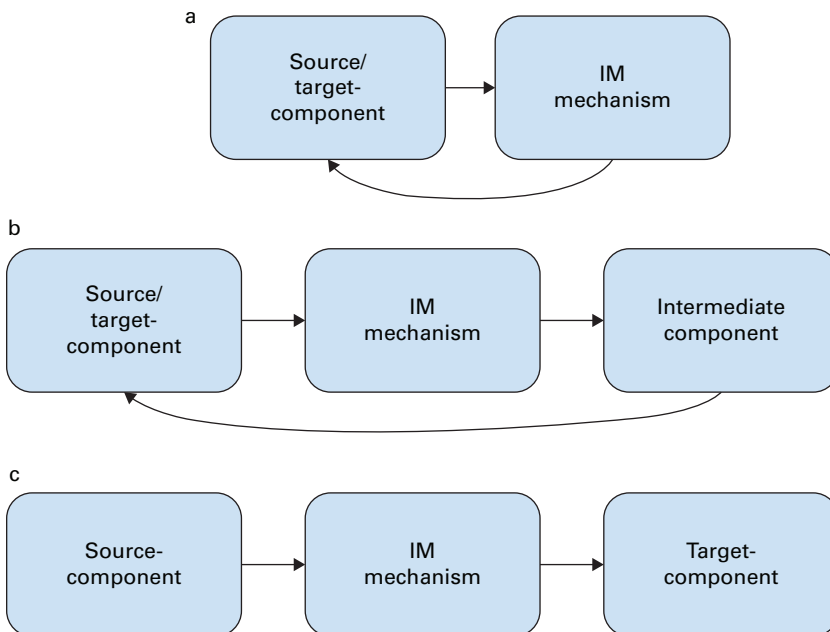


Figure 13.1 The key components of eIMs. (a) Case in which the source component and target component are the same structure. (b) Case in which the source component and the target component are the same structure, but the retroaction is mediated by an intermediate component. (c) Case in which the source component and the target component are different structures.

will be called *other IMs* (oIMs) to distinguish them from eIMs. Sometimes such “other mechanisms” mimic the acquisition of knowledge by a possible source component, but the latter is not actually present. For example, count-based novelty mechanisms (Bellemare et al. 2016) perform novelty detection on the basis of the frequency with which states are encountered rather than on the basis of how well they are memorized (although it is true that they are still present/absent in the counter memory). In other cases, other mechanisms are used that can support the function of acquiring knowledge and skills, but they themselves do not rely on a mechanism measuring the knowledge of some component. For example, the principle of empowerment (Klyubin et al. 2005), further discussed below, or the concept of bottlenecks (McGovern and Barto 2001), can support the acquisition of skills not by measuring the knowledge of a source component but by considering some properties of the environment or of the agent’s actions.

A critical difference between EMs and all IMs is the time when they express their function—that is, their utility. EMs tend to express their function at a time very close to when they are triggered. This is because they lead to the acquisition and consumption of material resources (organisms) or to the accomplishment of the user’s goals (robots), and when this happens they manifest their utility. Instead, IMs lead to the acquisition of knowledge and skills that are useful only later with respect to the time when they operate: the utility is indeed expressed only when such knowledge and skills are used to accomplish material resources or solve the user’s goals.

The time when IMs and EMs express their utility is particularly important because it makes it difficult to actually measure the effectiveness of a given IM mechanism. A possible way to measure such effectiveness is to divide the life of the agent into two phases (Schembri et al. 2007; Baldassarre et al. 2019): 1) the *intrinsic motivation phase*, in which the agent uses IMs to acquire knowledge and skills without a direct utility; 2) the *extrinsic motivation phase*, in which the agent uses the knowledge and skills acquired in the intrinsic phase to solve extrinsic problems. These two phases resemble the two main phases of human life involving a first infancy/childhood phase, mainly guided by IMs, and an adulthood phase, mainly guided by EMs (Schembri et al. 2007). This idea of the two phases was set at the core of the REAL competition (Robot open-Ended Autonomous Learning; Cartoni et al. 2020) proposed to create a benchmark for open-ended learning. In this competition, during a first intrinsic phase a simulated camera-arm-gripper robot can freely interact with some objects to autonomously acquire knowledge and skills without being given any goal or reward; in a second extrinsic phase, the quality of such knowledge and skills is measured by asking the robot to solve some sampled extrinsic tasks involving the re-creation of some sampled object configurations. The robot’s performance in the second phase thus furnishes a measure of the quality of the IM mechanisms used to acquire the knowledge in the first intrinsic phase. Two caveats come with this issue. Often in organisms, but also robots, IMs and EMs operate at the same time; for example, a robot might aim to learn how to manipulate an object while accomplishing a user’s tasks. This requires suitable arbitration mechanisms to mediate between the time and resources dedicated to IMs and EMs. Second, IM and EM mechanisms and functions might be mixed. For example, a “source component” and an “IM mechanism” might support a “target component” pursuing an extrinsic goal. For example, the next sections show a common use of novelty-based IMs to improve exploration in the accomplishment of extrinsic RL tasks.

EMs and eIMs (and sometimes also oIMs) also have a typical “temporal signature” (Bal-dassarre 2011). In particular, EMs tend to go away when the resources they are directed at are obtained and to come back when such resources are consumed/lost. For example, hunger and the reward of food ingestion go away after a sufficient amount of food is ingested and, say, blood glucose level increases and come back when the blood glucose level is low again. Instead, eIMs triggered by the acquisition of a particular piece of information stored in the source component tend to go away forever when such a piece of information is acquired (unless the information is forgotten). From a cognitive perspective, this helps in recognizing whether a motivation is an (e)IM or an EM; from a computational perspective, this is relevant because it possibly generates nonstationary, challenging problems (e.g., a typical problem faced is that if an IM mechanism is used to produce a reward signal for an RL component, then the resulting reward function keeps changing and so should the behavior).

13.2.2 Three Classes of eIMs

The computational literature has greatly contributed to distinguishing between different classes of IM mechanisms. These classes in particular involve eIMs and often are not applicable to oIMs: the classification presented here uses the term “IMs” to stay with the common nomenclature, but it actually refers to eIMs. A first contribution (cf. Oudeyer and Kaplan 2007) distinguishes between knowledge-based IMs, related to the acquisition of information on the world, and competence-based IMs (CB-IMs), related to the acquisition of the capacity to act effectively. Another contribution (Barto et al. 2013) highlights the need to differentiate between two types of knowledge-based IMs—namely, novelty-based IMs (NB-IMs) and prediction-based IMs (PB-IMs), often confused within the computational and biological/cognitive literature. The main features of these three classes of IMs, summarized in table 13.2, are now considered in detail. The classes are based on the function

Table 13.2
The three classes of (e)IMs

	Novelty-based IMs	Prediction-based IMs	Competence-based IMs
Source component: nature	Memory component (pattern magazine)	Predictor (forward model)	Skill (inverse model)
Source component: function	Pattern storing and recoding	Prediction of patterns based on other patterns	Action selection
IM mechanism: type of knowledge measured	How well represented is the item in memory, or how much did its representation improve?	What is the prediction error or the prediction error change?	How efficient/effective is the skill to accomplish the task/goal?
IM mechanism: processes involved in the measurement	One process: memory check	Two processes: (a) prediction (b) comparison of prediction with data	Multiple processes: iterated perception-action performance, check of success
Target component: typical functions	- Store/recode new items - Direct attention - Drive physical exploration - Support goal formation	- Improve predictions - Drive physical exploration - Direct attention - Support goal formation	- Speed up the learning of multiple skills

Source: Partially based on Barto et al. 2013.

implemented by the source component. For each class, there exist many subclasses depending on the functions and mechanisms of the target component. The IM mechanism always measures the level or change of the knowledge of the source component.

NB-IMs are based on a memory source component that encodes patterns, such as percepts, with the function of storing and possibly recoding them in more useful formats—for example, to compress information or to facilitate downstream processes. The IM mechanism of NB-IMs measures knowledge of the source component based on a one-step process that checks the level of novelty/familiarity of a target pattern, such as an image from the world. Another possibility is that the IM mechanism measures the novelty change of the internal representation of the pattern, rather than its level: this can happen if the pattern is experienced multiple times and the source component progressively improves its representation. Typical functions realized by the target component involve storing/recoding novel items (which is the case when the source and target components coincide), directing attention to novel items, driving their physical exploration, or supporting goal formation.

PB-IMs are based on a predictor source component that predicts patterns on the basis of other patterns. In particular, the predictor receives as input a pattern, and possibly the agent's action, and on this basis predicts a target pattern in a future time. The “future time” involves a time range in which the target item should happen, but predictions can also be “in space,” as in this example: “Given that I see a tree, I predict to see an apple if I look down 1 m.” The IM mechanism of PB-IMs performs a measurement of the knowledge of the source component (predictor) based on a two-step process in which first the predictor predicts the target pattern on the basis of an input pattern, and possibly of the agent's action, and then the mechanism compares the prediction with the actual target pattern to compute the size of their mismatch—that is, to compute the prediction error. Another possibility is that the measure involves the prediction error improvement (change), rather than the prediction error (level), based on monitoring how the prediction error evolves in time. Typical functions played by the target component, possibly coincident with the source component, involve improving predictions, directing attention to unpredicted items, driving their physical exploration, and forming goals.

CB-IMs assume the existence of tasks/goals and are based on a skill source component that can accomplish the tasks/goals (e.g., within a given period of time, the trial). The skill is a closed-loop or open-loop controller (e.g., a dynamic movement primitive, a policy, or an option) potentially able to solve the task/achieve the goal. The IM mechanism of CB-IMs performs a measurement of the knowledge of the source component that involves a multistep process: 1) the skill acts to accomplish the task/goal, possibly based on multiple sensorimotor steps; 2) its competence level is measured, for example, in terms of the amount of reward collected during the trial, or in terms of goal achievement, or in terms of distance between the achieved state and the goal. Another possibility is that the IM mechanism actually measures the competence improvement, rather than the competence level, based on the monitoring of the performance at multiple times. CB-IMs are particularly important in cases where multiple skills for accomplishing different tasks/goals have to be learned. In this respect, typical functions realized by the target component, usually coincident with the source component, are to learn multiple skills/goals, and the IM mechanism speeds up their learning by focusing on the skills with the highest learning speed.

Note that the definition of CB-IMs assumes the existence of tasks/goals. This is a critical aspect of CB-IMs because open-ended learning agents should be able to autonomously generate or discover such tasks/goals, as these are a major means to learn skills in an incremental fashion (Mirolli and Baldassarre 2013). Various oIMs considered in the following sections can be used to support such self-generation/discovery of tasks/goals.

13.3 Cognitive Robotics and Machine-Learning Models

This section considers the main functions that can be supported by IMs through the presentation of some computational models drawn from the robotics and machine-learning literature. In particular, it focuses on how IMs serve the acquisition of the overall capacity of agents to interact in the world to modify it (Mirolli and Baldassarre 2013). This focus leads us to consider in particular the relation between IMs and RL, the learning paradigm most closely related to the acquisition of the capacity to act in the world. Given this focus, the IM functions considered here are as follows (figure 13.2): (a) the accomplishment of sparse extrinsic rewards; (b) the self-generation of goals; (c) the acquisition of skills, either as policies per se or as policies linked to goals. These functions in particular are accomplished through processes that rely strongly on IM mechanisms alongside other mechanisms; these other mechanisms are 1) exploration, 2) goal sampling, imagination, or “marking,” and 3) the autonomous selection of skills to learn. Evolutionary processes are also considered to be possible general mechanisms searching the IM mechanisms themselves or the goals supporting CB-IMs.

13.3.1 Sparse Rewards

A first main function of IMs is to support the solution of RL tasks involving *sparse* extrinsic rewards—that is, rewards that are encountered rarely if the agent explores the environment randomly. Sparse rewards challenge learning agents, as they can be experienced only

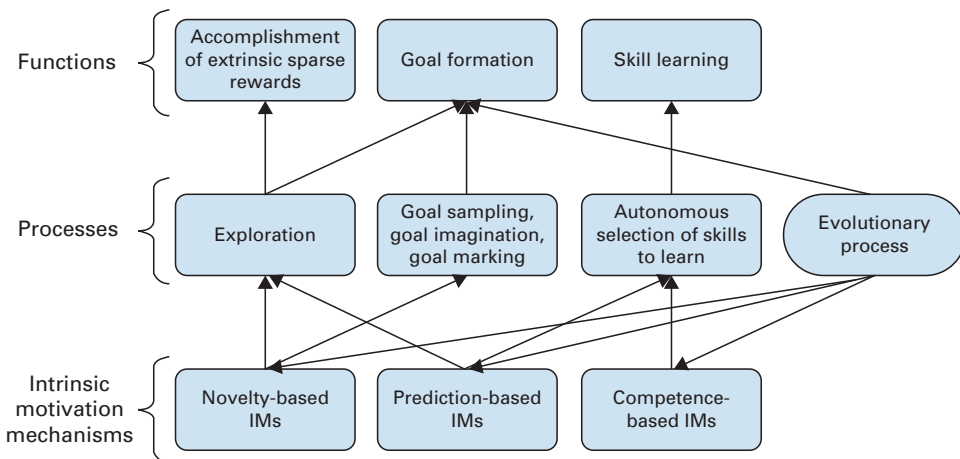


Figure 13.2 Some important functions that can be accomplished through IM mechanisms via some relevant processes.

after the performance of a long sequence of actions and therefore provide only very weak guidance for training. For example, imagine a camera-arm robot with no initial motor skills getting rewarded only for succeeding to grasp and lift an object with random movements. In this case it is close to impossible for a random exploration to lead to getting the reward and support learning. IMs can be very useful to solve tasks involving sparse rewards because they can facilitate the exploration of the environment through which the agent searches for the reward. Standard exploration methods, such as ϵ -greedy exploration (the agent selects a random action with a probability ϵ and the best action otherwise) and the Boltzmann distribution exploration (the possible actions are selected on the basis of a softmax function of their expected reward returns), are not adequate to face sparse-reward tasks because they lead to obtaining the reward only rarely. Various approaches have been proposed to produce a more effective exploration of the environment. A popular approach to foster exploration is based on NB-IMs. The idea is that the agent is attracted to states that it visited few times and tends to move away from familiar states. An extra reward (*novelty bonus*) could be given to the agent for making novel states attractive (Brafman and Tenenbholz 2002; Kakade and Dayan 2002). A nice property of novelty bonuses, and in general of IMs used to foster the pursuit of extrinsic rewards, is that since IMs have a transient nature, they tend to not affect the final policy acquired to maximize the final extrinsic reward.

A relevant class of methods using novelty to foster exploration in the search for extrinsic rewards is based on state novelty, measured as the number of times a state is encountered (Bellemare et al. 2016). In particular, these methods use density models to compute a pseudo-count of the times in which states are visited based on the generalization of the counts for similar states. The method was successfully applied to agents able to solve the Atari game *Montezuma's Revenge*, involving a highly sparse reward. Another model used for similar purposes is presented in Burda et al. (2018). Here a random network is used to recode the state observations (images), and a second “copy” network is trained with supervised learning to “mimic” the first network (same input; desired output as the random network). The idea is that when states become more familiar the error of the copy network decreases.

Exploration to pursue extrinsic goals could also be pursued through PB-IMs. PB-IMs can rely on the prediction error (Schmidhuber 1991b) or the prediction error improvement (Schmidhuber 1991a) of a predictor network—that is, a world model predicting the next state on the basis of the current state and possibly the planned action. The prediction error has the disadvantage, if used by an IM mechanism, of not fading away in stochastic worlds. This problem is solved by the prediction error improvement, although at the cost of having a noisy and slow-adjusting signal. In the initial models using these strategies (Schmidhuber 1991a, 1991b), the predictor was used both as the source component and as the target component, meaning that the function of the used IM was to train the predictor itself. The same IM mechanism can, however, be also used to foster exploration to accomplish extrinsic tasks involving sparse rewards. An example of a model doing this is presented in Pathak et al. (2017). Here a forward model is used to produce a prediction error used as an intrinsic reward to train a RL agent to solve video games, such as *Mario Bros.*, involving sparse extrinsic rewards. Interestingly, the model also proposes a mechanism to only focus on effects that are caused by the agent's actions by using a predictor that uses as input the internal representations of an inverse model predicting actions based on an input formed by the before-action state and the after-action state.

13.3.2 Goal Formation

A very interesting function for which IMs can be used is related to the acquisition of multiple sensorimotor skills that might be later used to accomplish other intrinsic tasks, or extrinsic tasks, particularly within a hierarchical RL framework where behavior is chunked into options (Sutton et al. 1999). Here we consider the goal-based version of options, in which each option involves (Barto et al. 2004; Singh et al. 2004) 1) a termination condition associated with the accomplishment of a goal, 2) an action policy indicating the primitive actions to select in correspondence with different states of the world, 3) possibly an initiation set encompassing the states from which if executed the policy is able to accomplish the goal. A goal is a representation of a set of world states that if reactivated internally drives the agent to act in the world so that the world assumes one of those states. There are various types of goals, such as goals as states of the world, goals as trajectories of states, avoidance goals, maintenance goals, and more (Merrick et al. 2016), but here we focus only on state goals for simplicity, and as many considerations can be extended to the other types of goals. Goals can have different levels of abstraction and can involve one's own body (Mannella et al. 2018; Hoffmann et al. 2010), the external environment (e.g., Santucci et al. 2016), the relation between a couple of elements (Kulkarni et al. 2016), or social aspects (Acevedo-Valle et al. 2018).

Various subfunctions, supported by IMs, are important for learning repertoires of multiple skills for later use. Here four are considered: 1) the autonomous generation of goals; 2) the coverage of the widest possible part of the goal space (goal exploration); 3) the generation of the reward for learning the policy of the single option; 4) the support of the progressive learning of skills, from easy to difficult, to speed up their acquisition.

The function of goal formation is important because during intrinsic learning the robots are not given any task to solve and so should autonomously self-generate tasks/goals guiding the acquisition of the related skills. Note that although goal formation is extremely important for open-ended learning, and various methods supporting it involve eIMs (Mirolli and Baldassarre 2013), it often also involves mechanisms differing from the elements of eIMs (source component, IM mechanism, and target component). These are here considered oIM mechanisms; further investigations are needed to understand if and how oIMs are linked to eIMs. We will now consider some relevant methods used to autonomously generate goals.

Goal sampling

When the goal space is given—for example, it is formed by the posture angles of a robot or the x, y positions of an object on a table—goals can be sampled on the basis of their skill learnability. For example, *goal babbling* (Rolf et al. 2010) allows a robot to self-generate posture goals that facilitate the learning of a coherent inverse model by maximizing the end-effector displacement, which favors the exploration of novel goals while minimizing the posture change, which favors the learning of regular versus awkward postures among the possible redundant postures. The approach has been later extended, for example, to learn multiple models in parallel (from end-effector position space to joint space and from the joint space to the motor space) through associative radial-basis-function networks growing on the basis of novel experiences (Rayyes and Steil 2019).

The goal space might not be given to the agent but form a subspace of the state or observation space to be actively searched. In this case goal sampling is not possible,

especially if the subspace is small with respect to the whole space; in this case the goal subspace has to be actively discovered by the agent. Consider, for example, an observation space formed by images. In this case, the agent has to actively discover the image goals that it might actually achieve with its actions within the whole huge space formed by all possible images corresponding to all combinations of the pixel values. Now some approaches usable to this purpose are considered.

Goal marking

A number of models have proposed specific mechanisms to “mark”—that is, establish as goals—experienced states or observations. These models do not have the features of eIMs but can support open-ended learning via the formation of goals and the learning of the related skills, so they can be considered oIMs. A classic approach is the one for marking as goals the experienced states of the world that represent bottlenecks (McGovern and Barto 2001), nodal conditions that are often traversed when solving multiple extrinsic tasks (e.g., doorways when navigating an office).

Another model proposed to form goals corresponding to salient events, such as a change of light or sound (Barto et al. 2004; Singh et al. 2004). Linked to this, another approach proposed to mark as goals the novel observations that follow changes caused by the agent’s actions in the environment (Santucci et al. 2016; Mannella et al. 2018). The idea behind this approach is that what robots (and organisms) ultimately should do during intrinsic learning is become able to change the world at will, so the observations that follow a change caused by own actions indicate a potential for doing this. The novelty of the changes guarantees that the goal has not been already formed. If changes in the world can also happen independently of the agent’s action, additional mechanisms are needed to allow the agent to identify the subset of changes that depend on its action (Sperati and Baldassarre 2018; Pathak et al. 2017). Another approach forms goals when a particular relation between couples of elements takes place—for example, the “agent” picks up a “key” in an Atari game (Kulkarni et al. 2016).

A different approach (Zhao et al. 2012) uses RL to acquire various behaviors with motorized cameras within an active vision context (Ballard 1991; Ognibene and Baldassarre 2015)—for example, to lead two cameras to focus on the same target (vergence control). Here the model uses as a reward the accuracy of the reconstruction of images of a sparse-coding component (Olshausen and Field 1996), and the low error marks states where the two cameras manage to focus on the same target.

Another approach for skill learning is empowerment (Klyubin et al. 2005). Empowerment has a wide relevance for open-ended learning, but for lack of space only a few elements of it can be considered here. Empowerment is based on information theory and can be used to assign to each given world state a value that represents the variety of different outcome states that the agent can achieve with its actions from the given state. States with high empowerment can be used as target states; for example, their empowerment value can be directly used as reward to drive skill learning (T. Jung et al. 2011). Der and Martius (2015) propose another approach exploiting emergent properties of the environment-body-controller dynamics to autonomously acquire interesting motor skills in dynamic simulated agents. The skills are acquired on the basis of a simple two-layer neural network sensorimotor controller whose connection weights are trained through a differential extrinsic

plasticity (DEP) rule derived from differential Hebbian learning (Zappacosta et al. 2018) that captures correlations between the changes of the input neurons and the output neurons.

Goal manifold search

This strategy searches goals within large observation spaces based on the idea that similar goals involve similar skills/actions, and so the performance of noisy variants of the already discovered skills/actions might possibly lead to discovering new achievable goals. This strategy was first used in a model (*skill babbling*; Reinhart 2017) to control an arm robot learning to displace an object in the 3D space. The model forms clusters of similar goals and discovers new goals by performing noisy versions of the actions corresponding to the centroid goals of clusters. The active goal manifold exploration model (AGME; Cartoni and Baldassarre 2018) actively discovers the goal manifold hidden in the observation space—for example, a posture space or an image space. For this purpose, the model builds a distance-based graph of the discovered goals, selects goals that have a higher distance from other discovered goals, generates perturbed versions of the policies associated with such goals, and performs them to discover new goals. The quality diversity algorithm (Kim et al. 2019) learns a repertoire of behaviors and goals by searching for behaviors that are different (novel) with respect to the already learned behaviors. The algorithm is, for example, used to allow a humanoid robot to acquire the skills to throw a ball into a basket located in many possible different positions (goals) on the floor. The hindsight experience replay approach (HER; Adrychowicz et al. 2017) exploits the outcome of policies to discover new goals, even if they are different from the pursued goal. The approach is very effective to incrementally discover new goals—for example, to manipulate objects in a simulated camera-arm-gripper robot.

Goal formation by imagination

Another related strategy discovers goals by first imagining them. For example, the reinforcement learning with imagined goals model (RIG; Nair et al. 2018), tested with a robot arm moving objects on a table, uses a generative model (a variational autoencoder; Kingma and Welling 2013) to first learn an internal compact representation of goals by randomly exploring the environment and then to “imagine” other possible goals whose skills are learned by RL. A later version of the model generates goals that have a high probability of being novel with respect to already learned goals by sampling them on the fringe of the distribution of the internal representation of the discovered goals (Pong et al. 2019). “Imagination” is a relevant means not only to generate goals but also to formulate plans to achieve those goals by assembling other goals/skills (Seepanomwan et al. 2015; Hung et al. 2018; M. Jung et al. 2019; Tanneberg et al. 2019) possibly acquired with IMs. This is an interesting trend that reformulates some high-level concepts elaborated by the classic symbolic planning literature (Russell and Norvig 2016), such as goals and planning, through neural network representations.

13.3.3 Selection of Skills to Train

The literature on animal learning (Skinner 1953) and on staged child development (Piaget 1953) shows that learning progress is faster if it proceeds from easy to difficult tasks. This strategy can also be used in artificial systems by training them with a curriculum involving increasingly difficult tasks (Asada et al. 1996; Bengio et al. 2009). One of the most interesting

uses of IMs allows open-ended learning agents to autonomously select the skills needed to train to achieve goals possibly generated autonomously with the approaches illustrated above. Initially, PB-IMs were used to support the autonomous selection of tasks to learn (e.g., Singh et al. 2004; Oudeyer et al. 2007). Here the source component was a predictor, while the target component was the skill to learn, and the agent focused learning on skills causing the highest predictor error, or prediction error improvement, of the predicted skill outcome. Successively, CB-IMs were shown to be more appropriate than PB-IMs for selecting the skills to train because the predictor of the PB-IMs might learn to predict the skill outcome too early or too late with respect to when the controller finishes learning the skill. Instead, CB-IMs directly measure the competence acquired by different skills so it returns accurate information usable for selecting them (Santucci et al. [2013] compared these different IM mechanisms for task selection).

When a goal can be accomplished starting from a different initial condition, the CB-IM signal related to the goal must also take into account such an initial condition; moreover, when a goal can be selected not only depending on its learning rate but also depending on whether its achievement can be the precondition for learning other skills, then the CB-IM signal has to be used as a reward within a whole RL process selecting goals rather than actions (Santucci et al. 2019). IMs can also guide the progressive learning of increasingly difficult tasks represented at multiple levels of abstraction—for example, in robots learning to interact with different objects (Ugur and Piater 2016). In all these models, the skill of the selected goal should be trained (with RL) through a pseudo-reward equal to one when the goal is accomplished and to zero otherwise. This is more effective than what was done in the early years of research on IMs when the PB-IM signal used to select the goal/skill was also used to train the skill, as the PB-IM signal gradually fades away when the skill is learned.

13.3.4 Evolution

Tasks/goals could also be generated autonomously through evolutionary processes (genetic algorithms). Schembri et al. (2007) proposed the first model to do so in a population of RL simulated robots moving on a colored arena. During the intrinsic phase, the robots used intrinsic reward functions generated by a genetic algorithm to learn skills. In the later extrinsic phase, the robots learned to compose the acquired skills to accomplish extrinsic tasks (specific places in the arena). The success in learning these extrinsic tasks produced the fitness for the genetic algorithm. Singh et al. (2010) used an algorithm equivalent to evolution to search reward functions of RL agents engaged in searching for food in a grid world. They found that reward functions having the highest score rewarded the agents not only for searching for food but also for “opening boxes” where food was hidden. The model was used to suggest the existence of a continuum between EMs and IMs, rather than a distinction between them, as from an evolutionary perspective the two differ only for their distance from the events increasing fitness. The view proposed here distinguishes eIMs and EMs, as eIMs are based on the measure of *knowledge in a component of the controller*, whereas EMs are based on the measure of *material resources in the body or the environment*. It is, however, true that in the case of evolved oIMs that support the formation of goals and skills, as in the models reviewed above, a continuum with EMs can be seen since the criterion of the “knowledge-measurement” typical of eIMs is missing.

There is an additional important problem for open-ended learning that could be tackled with evolutionary approaches: Which goals/skills should be acquired, among those possible, to later best learn several different extrinsic tasks in a given domain? Del Verme et al. (2020) faced this problem and used a genetic algorithm to search goals/skills that were optimal for the solution of tasks drawn from a certain distribution of possible tasks in a given environment. The work showed how the optimal goals and skills depended on the time budget that the agent had in order to solve the extrinsic tasks and on the physical regularities of the environment. It so demonstrated that “fixed” mechanisms for goal generation, as those seen above, might lead to suboptimal solutions. Importantly, evolutionary approaches might thus be used to evolve the IM mechanisms themselves, as hinted by the arrows in figure 13.2 departing from the “evolutionary processes” box (Salgado et al. 2016). Although very interesting, this possibility is now limited by its high computational costs.

13.4 Conclusion

The study of intrinsic motivations is making important progress. However, many relevant open issues need further investigation. One open issue is the clarification of how non-epistemic intrinsic motivations work and are related to epistemic ones. Another open issue is the clarification of the link between intrinsic motivations and the autonomous formation of goals. A further issue, in part related to that, is the clarification of the relationship existing between intrinsic motivations and concepts such as empowerment and sensorimotor emergent behaviors. We have also seen how the computational literature is uncovering the existence of an articulated typology of intrinsic motivation mechanisms and functions. Understanding if and how these are also present in organisms’ brains and behavior is a very interesting open problem.

Robot open-ended learning itself is still unsolved, as shown by the fact that we do not have robots able to undergo a truly open-ended learning experience leading to an unbounded accumulation of knowledge and skills. This might depend on multiple factors. On the side of goal formation, we have various mechanisms for the autonomous generation of goals, but all of them have limitations: goal sampling can only be applied to known small goal spaces; goal formation based on mechanisms such as bottlenecks, novel environment changes, goal-manifold discovery, and goal imagination has yet to be scaled to larger goal spaces and different domains. The autonomous selection of skills to train, based on competence-based intrinsic motivations, is becoming a standard, but it generally assumes discrete goals and hence must be further developed to be easily applicable to continuous goal spaces. Finally, systems working with discrete goals solve extrinsic problems based on planning and search methods that require the number of learned goal/skills to be limited to be efficient. This problem might be solved with evolutionary methods that indirectly search for a few robust skills to learn by searching the IM mechanisms themselves that lead to their generation; this, however, currently has a prohibitive computational cost.

Despite these challenges, the research field of open-ended learning driven by intrinsic motivations is surely one of the most exciting fields of cognitive robotics due to its potential for applications in robots acting in unstructured environments and to its close link with some of the most sophisticated and intriguing processes of human cognition, such as curiosity and the drive for the autonomous acquisition of knowledge.

Acknowledgments

This research received funding from the European Union under the 7th Research Framework Program, Grant Agreement n° 231722, Project “IM-CLeVeR—Intrinsically Motivated Cumulative Learning Versatile Robots”; and the Horizon 2020 Research and Innovation Program, Grant Agreement n° 713010, Project “GOAL-Robots—Goal-based Open-ended Autonomous Learning Robots.”

Additional Reading and Resources

- A collection of works on intrinsic motivations and open-ended learning: Baldassarre, Gianluca, and Marco Mirolli. 2013. *Intrinsically Motivated Learning in Natural and Artificial Systems*. Berlin: Springer.
- A work that complements the current work, with a perspective on the biology and brain mechanisms underlying intrinsic motivations: Baldassarre, Gianluca. 2011. “What Are Intrinsic Motivations? A Biological Perspective.” In *Proceedings of the International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob-2011)*. New York: IEEE.
- A work presenting a general architecture supporting several of the functions for open-ended learning discussed in the chapter and usable for domains involving discrete goals: Santucci, Vieri Giuliano, Gianluca Baldassarre, and Marco Mirolli. 2016. “GRAIL: A Goal-Discovering Robotic Architecture for Intrinsically-Motivated Learning.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 214–231. doi:10.1109/tcds.2016.2538961.
- Link to the project that sponsored this research, which furnishes additional resources and software on the topic: www.goal-robots.eu.

References

- Acevedo-Valle, Juan M., Verena V. Hafner, and Cecilio Angulo. 2018. “Social Reinforcement in Artificial Pre-linguistic Development: A Study Using Intrinsically Motivated Exploration Architectures.” *IEEE Transactions on Cognitive and Developmental Systems* 12 (2): 198–208.
- Andrychowicz, Marcin, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. 2017. “Hindsight Experience Replay.” ArXiv preprint: 1707.01495v3.
- Asada, Minoru, Shoichi Noda, Sukoya Tawaratsumida, and Koh Hosoda. 1996. “Purposive Behavior Acquisition for a Real Robot by Vision-Based Reinforcement Learning.” *Machine Learning* 23 (2–3): 279–303.
- Baldassarre, Gianluca. 2011. “What Are Intrinsic Motivations? A Biological Perspective.” In Vol. 2, *2011 IEEE International Conference on Development and Learning*, 1–8. New York: IEEE.
- Baldassarre, Gianluca, William Lord, Giovanni Granato, and Vieri Giuliano Santucci. 2019. “An Embodied Agent Learning Affordances with Intrinsic Motivations and Solving Extrinsic Tasks with Attention and One-Step Planning.” *Frontiers in Neurobotics* 13 (45): e1–e26.
- Baldassarre, Gianluca, and Marco Mirolli. 2013. *Intrinsically Motivated Learning in Natural and Artificial Systems*. Berlin: Springer.
- Ballard, Dana H. 1991. “Animate Vision.” *Artificial Intelligence* 48:57–86.
- Barto, Andrew, Marco Mirolli, and Gianluca Baldassarre. 2013. “Novelty or Surprise?” *Frontiers in Psychology* 4 (907): e1–e15.
- Barto, Andrew G., Satinder Singh, and Nuttpong Chentanez. 2004. “Intrinsically Motivated Learning of Hierarchical Collections of Skills.” In *Proceedings of the 3rd International Conference on Developmental Learning (ICDL2004)*, 112–119. New York: IEEE.

- Bellemare, Marc G., Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. "Unifying Count-Based Exploration and Intrinsic Motivation." ArXiv preprint: 1606.01868.
- Bengio, Yoshua, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. "Curriculum Learning." In *Proceedings of the 26th Annual International Conference on Machine Learning*, 41–48. New York: Association for Computing Machinery.
- Berlyne, Daniel E. 1966. "Curiosity and Exploration." *Science* 153 (3731): 25–33.
- Brafman, Ronen I., and Moshe Tennenholtz. 2002. "R-Max—a General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning." *Journal of Machine Learning Research* 3:213–231.
- Burda, Yuri, Harrison Edwards, Amos Storkey, and Oleg Klimov. 2018. "Exploration by Random Network Distillation." ArXiv preprint: 1810.12894v1.
- Cartoni, Emilio, and Gianluca Baldassarre. 2018. "Autonomous Discovery of the Goal Space to Learn a Parameterized Skill." ArXiv preprint: 1805.07547v1.
- Cartoni, Emilio, Francesco Mannella, Vieri Giuliano Santucci, Jochen Triesch, Elmar Rueckert, and Gianluca Baldassarre. 2020. "REAL-2019: Robot Open-Ended Autonomous Learning Competition." In *Proceedings of Machine Learning Research* 1:142–152.
- Del Verme, Manuel, Bruno Castro Da Silva, and Gianluca Baldassarre. 2020. "Optimal Options for Multi-task Reinforcement Learning under Time Constraints." ArXiv preprint: 2001.01620v1.
- Der, Ralf, and Georg Martius. 2015. "Novel Plasticity Rule Can Explain the Development of Sensorimotor Intelligence." *Proceedings of the National Academy of Science USA* 112 (45): e6224–e6232.
- Doncieux, Stéphane, David Filliat, Natalia Díaz-Rodríguez, Timothy Hospedales, Richard Duro, Alexandre Coninx, Diederik M. Roijers, Benoît Girard, Nicolas Perrin, and Olivier Sigaud. 2018. "Open-Ended Learning: A Conceptual Framework Based on Representational Redescription." *Frontiers in Neurobotics* 12 (59): e1–e6.
- Doya, Kenji, and Tadahiro Taniguchi. 2019. "Toward Evolutionary and Developmental Intelligence." *Current Opinion in Behavioral Sciences* 29:91–96.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2017. *Deep Learning*. Cambridge, MA: MIT Press.
- Harlow, Harry F. 1950. "Learning and Satiation of Response in Intrinsically Motivated Complex Puzzle Performance by Monkeys." *Journal of Comparative and Physiological Psychology* 43 (4): 289–294.
- Hoffmann, Matej, Hugo Marques, Alejandro Arieta, Hidenobu Sumioka, Max Lungarella, and Rolf Pfeifer. 2010. "Body Schema in Robotics: A Review." *IEEE Transactions on Autonomous Mental Development* 2 (4): 304–324.
- Hull, Clark L. 1943. *Principles of Behavior*. New York: Appleton-Century-Crofts.
- Hung, Chia-Chun, Timothy Lillicrap, Josh Abramson, Yan Wu, Mehdi Mirza, Federico Carnevale, Arun Ahuja, and Greg Wayne. 2018. "Optimizing Agent Behavior over Long Time Scales by Transporting Value." ArXiv preprint: 1810.06721v1.
- Jacquey, Lisa, Gianluca Baldassarre, Vieri Giuliano Santucci, and J. Kevin O'Regan. 2019. "Sensorimotor Continuities as a Key Drive of Development: From Babies to Robots." *Frontiers in Neurobotics* 13 (98): e1–e20.
- Jung, Minju, Takazumi Matsumoto, and Jun Tani. 2019. "Goal-Directed Behavior under Variational Predictive Coding: Dynamic Organization of Visual Attention and Working Memory." ArXiv preprint: 1903.04932v1.
- Jung, Tani, Daniel Polani, and Peter Stone. 2011. "Empowerment for Continuous Agent Environment Systems." *Adaptive Behavior* 19 (1): 16–39.
- Kakade, Sham, and Peter Dayan. 2002. "Dopamine: Generalization and Bonuses." *Neural Networks* 15 (4–6): 549–559.
- Kim, Seungsu, Alexandre Coninx, and Stéphane Doncieux. 2019. "From Exploration to Control: Learning Object Manipulation Skills through Novelty Search and Local Adaptation." ArXiv preprint: 1901.00811v1.
- Kingma, Diederik P., and Max Welling. 2013. "Auto-Encoding Variational Bayes." ArXiv preprint: 1312.6114v10.
- Kish, George Bela. 1955. "Learning When the Onset of Illumination Is Used as the Reinforcing Stimulus." *Journal of Comparative and Physiological Psychology* 48 (4): 261–264.
- Klyubin, Alexander S., Daniel Polani, and Christopher L. Nehaniv. 2005. "All Else Being Equal Be Empowered." In *European Conference on Artificial Life*, 744–753. Berlin: Springer.
- Kulkarni, Tejas D., Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. 2016. "Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation." In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 3682–3690. New York: Curran Associates.
- Lungarella, Max, Giorgio Metta, Rolf Pfeifer, and Giulio Sandini. 2003. "Developmental Robotics: A Survey." *Connection Science* 15 (4): 151–190.
- Mannella, Francesco, Vieri Giuliano Santucci, Somogyi Eszter, Lisa Jacquey, Kevin J. O'Regan, and Gianluca Baldassarre. 2018. "Know Your Body through Intrinsic Goals." *Frontiers in Neurobotics* 12:30.

- Marr, David, and Tomaso Poggio. 1976. "From Understanding Computation to Understanding Neural Circuitry." *A.I. Memo AIM-357*. Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.
- Marraffa, Rodolfo, Valerio Sperati, Daniele Caligiore, Jochen Triesch, and Gianluca Baldassarre. 2012. "A Bio-inspired Attention Model of Anticipation in Gaze-Contingency Experiments with Infants." In *2012 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics*, 1–8. New York: IEEE.
- McGovern, Amy, and Andrew G. Barto. 2001. "Automatic Discovery of Subgoals in Reinforcement Learning Using Diverse Density." Faculty Publication Series. University of Massachusetts Amherst, Computer Science Department.
- Merrick, Kathryn, Nazmul Siddique, and Inaki Rano. 2016. "Experience-Based Generation of Maintenance and Achievement Goals on a Mobile Robot." *Journal of Behavioral Robotics* 7 (1): 67–84.
- Mirolli, Marco, and Gianluca Baldassarre. 2013. "Functions and Mechanisms of Intrinsic Motivations: The Knowledge versus Competence Distinction." In *Intrinsically Motivated Learning in Natural and Artificial Systems*, 49–72. Berlin: Springer.
- Nair, Ashvin, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. 2018. "Visual Reinforcement Learning with Imagined Goals." In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 9209–9220. New York: Curran Associates.
- Ognibene, Dimitri, and Gianluca Baldassarre. 2015. "Ecological Active Vision: Four Bio-inspired Principles to Integrate Bottom-Up and Adaptive Top-Down Attention Tested with a Simple Camera-Arm Robot." *IEEE Transactions on Autonomous Mental Development* 7 (1): 3–25.
- Olshausen, Bruno, and David J. Field. 1996. "Emergence of Simple-Cell Receptive Field Properties by Learning a Sparse Code for Natural Images." *Nature* 381 (6583): 607–609.
- Oudeyer, Pierre-Yves, and Frederic Kaplan. 2007. "What Is Intrinsic Motivation? A Typology of Computational Approaches." *Frontiers in Neurobotics* 1:6.
- Oudeyer, Pierre-Yves, Frédéric Kaplan, and Verena V. Hafner. 2007. "Intrinsic Motivation Systems for Autonomous Mental Development." *IEEE Transactions on Evolutionary Computation* 11 (2): 265–286.
- Panksepp, Jaak. 1998. *Affective Neuroscience: The Foundations of Human and Animal Emotions*. Oxford: Oxford University Press.
- Pathak, Deepak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. 2017. "Curiosity-Driven Exploration by Self-Supervised Prediction." In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 488–489. New York: IEEE.
- Piaget, Jean. 1953. *The Origins of Intelligence in Children*. London: Routledge and Kegan Paul.
- Polizzi di Sorrentino, Eugenia, Gloria Sabbatini, Valentina Truppa, Anna Bordonali, Fabrizio Taffoni, Domenico Formica, Gianluca Baldassarre, Marco Mirolli, and Visalberghi Guglielmelli Eugenio. 2014. "Exploration and Learning in Capuchin Monkeys (*Sapajus Spp.*): The Role of Actionoutcome Contingencies." *Animal Cognition* 17 (5): 1081–1088.
- Pong, Vitchyr H., Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. 2019. "Skew-Fit: State-Covering Self-Supervised Reinforcement Learning." ArXiv preprint: 1903.03698v2.
- Rayyes, Rania, and Jochen Steil. 2019. "Online Associative Multi-stage Goal Babbling toward Versatile Learning of Sensorimotor Skills." In *2019 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics*, 327–334. New York: IEEE.
- Reinhart, Felix. 2017. "Autonomous Exploration of Motor Skills by Skill Babbling." *Autonomous Robots* 41 (7): 1521–1537.
- Rolf, Matthias, Jochen J. Steil, and Michael Gienger. 2010. "Goal Babbling Permits Direct Learning of Inverse Kinematics." *IEEE Transactions on Autonomous Mental Development* 2 (3): 216–229.
- Russell, Stuart J., and Peter Norvig. 2016. *Artificial Intelligence: A Modern Approach*. Harlow, UK: Pearson Education.
- Ryan, Richard M., and Edward L. Deci. 2000. "Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being." *American Psychologist* 55 (1): 68–78.
- Salgado, Rodrigo, Abraham Prieto, Pilar Caamaño, Francisco Bellas, and Richard J. Duro. 2016. "Motiven: Motivational Engine with Sub-Goal Identification for Autonomous Robots." In *IEEE Congress on Evolutionary Computation*, 4887–4894. New York: IEEE.
- Santucci, Vieri Giuliano, Gianluca Baldassarre, and Emilio Cartoni. 2019. "Autonomous Reinforcement Learning of Multiple Interrelated Tasks." In *2019 Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*, 221–227. New York: IEEE.
- Santucci, Vieri Giuliano, Gianluca Baldassarre, and Marco Mirolli. 2013. "Which Is the Best Intrinsic Motivation Signal for Learning Multiple Skills?" *Frontiers in Neurobotics* 7 (22): e1–e14.

- Santucci, Vieri Giuliano, Gianluca Baldassarre, and Marco Mirolli. 2016. "GRAIL: A Goal-Discovering Robotic Architecture for Intrinsically-Motivated Learning." *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 214–231.
- Schembri, Massimiliano, Marco Mirolli, and Gianluca Baldassarre. 2007. "Evolving Childhood's Length and Learning Parameters in an Intrinsically Motivated Reinforcement Learning Robot." In *the 7th International Conference on Epigenetic Robotics (Epirob2007)*, 141–148. Lund, Sweden: Lund University Press.
- Schmidhuber, Juergen. 1991a. "Curious Model-Building Control Systems." In *International Joint Conference on Artificial Neural Networks*, 1458–1463. Piscataway, NJ: IEEE.
- Schmidhuber, Juergen. 1991b. "A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers." In *International Conference on Simulation of Adaptive Behavior: From Animals To Animals*, 222–227. Boston: MIT Press.
- Schmidhuber, Juergen. 2010. "Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990–2010)." *IEEE Transactions on Autonomous Mental Development* 2 (3): 230–247.
- Seepanomwan, Kristiana, Daniele Caligiore, Angelo Cangelosi, and Gianluca Baldassarre. 2015. "Generalization, Decision Making, and Embodiment Effects in Mental Rotation: A Neurobotic Architecture Tested with a Humanoid Robot." *Neural Networks* 72:31–47.
- Seepanomwan, Kristiana, Vieri Giuliano Santucci, and Gianluca Baldassarre. 2017. "Intrinsically Motivated Discovered Outcomes Boost User's Goals Achievement in a Humanoid Robot." In *2017 Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*, 178–183. New York: IEEE.
- Singh, Satinder, Andrew G. Barto, and Nuttapon Chentanez. 2004. "Intrinsically Motivated Reinforcement Learning." In *Advances in Neural Information Processing Systems*, 1281–1288. Boston: MIT Press.
- Singh, Satinder, Richard L. Lewis, Andrew G. Barto, and Jonathan Sorg. 2010. "Intrinsically Motivated Reinforcement Learning: An Evolutionary Perspective." *IEEE Transactions on Autonomous Mental Development* 2 (2): 70–82.
- Skinner, Burrhus F. 1938. *The Behavior of Organisms*. New York: Appleton-Century-Crofts.
- Skinner, Burrhus F. 1953. *Science and Human Behavior*. New York: Macmillan.
- Sperati, Valerio, and Gianluca Baldassarre. 2018. "A Bio-inspired Model Learning Visual Goals and Attention Skills through Contingencies and Intrinsic Motivations." *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 326–344.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Sutton, Richard S., Doina Precup, and Satinder Singh. 1999. "Between MDPS and Semi-MDPS: A Framework for Temporal Abstraction in Reinforcement Learning." *Artificial Intelligence* 112:181–211.
- Taffoni, Fabrizio, Eleonora Tamilia, Valentina Focaroli, Domenico Formica, Luca Ricci, Giovanni Di Pino, Gianluca Baldassarre, Marco Mirolli, Eugenio Guglielmelli, and Flavio Keller. 2014. "Development of Goal-Directed Action Selection Guided by Intrinsic Motivations: An Experiment with Children." *Experimental Brain Research* 232 (7): 2167–2177.
- Tanneberg, Daniel, Jan Peters, and Elmar Rueckert. 2019. "Intrinsic Motivation and Mental Replay Enable Efficient Online Adaptation in Stochastic Recurrent Networks." *Neural Networks* 109:67–80.
- Timbergen, Nikolaas. 1963. "On Aims and Methods of Ethology." *Zeitschrift für Tierpsychologie* 20 (4): 410–433.
- Ugur, Emre, and Justus Piater. 2016. "Emergent Structuring of Interdependent Affordance Learning Tasks Using Intrinsic Motivation and Empirical Feature Selection." *IEEE Transactions on Cognitive and Developmental Systems* 9 (4): 328–340.
- Vieira Neto, Hugo, and Ulrich Nehmzow. 2007. "Visual Novelty Detection with Automatic Scale Selection." *Robotics and Autonomous Systems* 55 (9): 693–701.
- Weng, Juyang, James McClelland, Alex Pentland, Olaf Sporns, Ida Stockman, Mriganka Sur, and Esther Thelen. 2001. "Autonomous Mental Development by Robots and Animals." *Science* 291:599–600.
- White, Ruth W. 1959. "Motivation Reconsidered: The Concept of Competence." *Psychological Review* 66:297–333.
- Zappacosta, Stefano, Francesco Mannella, Marco Mirolli, and Gianluca Baldassarre. 2018. "General Differential Hebbian Learning: Capturing Temporal Relations between Events in Neural Networks and the Brain." *PLoS Computational Biology* 14 (8): e1006227.
- Zhao, Yu, Constantin A. Rothkopf, Jochen Triesch, and Bertram E. Shi. 2012. "A Unified Model of the Joint Development of Disparity Selectivity and Vergence Control." In *2012 IEEE Joint International Conference on Development and Learning and Epigenetic Robotics*, 1–6. New York: IEEE.
- Zlatev, Jordan, and Christian Balkenius. 2001. "Introduction: Why Epigenetic Robotics?" In *Proceedings of the First International Workshop on Epigenetic Robotics*, 1–4. Lund, Sweden: Lund University Press.

14 Principles of Cognitive Vision

Yiannis Aloimonos and Giulio Sandini

14.1 Is Perception Only a Recovery Process?

The goal of computer vision has been to create three-dimensional (3D) descriptions of the scene in view and to recognize the aspects of the scene by assigning labels to the objects and actions that exist or are happening in the scene (Marr 1982). These labels would then be provided to symbolic reasoning systems—of the kind artificial intelligence develops—that would reason about the world. Psychologists have also written extensively on the cognitive impenetrability of visual perception, implying that the workings of visual perception are shielded from any cognition. Thus, visual perception is seen as a mechanistic black box that delivers labels through recognition, and nowadays this is achieved through machine-learning techniques that use gargantuan amounts of (mostly static) data.

Current practice, however, has suggested repeatedly that the path from pixels to symbols in a bottom-up manner is difficult, if not impossible. It certainly does not make explicit the causal link between the present (i.e., what is seen now) and the past. Therefore, vision cannot be used to anticipate the future course of events, which is the core of cognition. It appears that knowledge of some form comes into the process of visual perception quite early in the perceptual process. In the classical framework described above, there is only one place where perception and cognition meet. However, this is counterintuitive to our common-sense understanding of perception and thinking. Human behavior is active. Humans (and animals) continuously shift their gaze to different locations in their scene in view. They recognize objects, sounds, and actions, and this leads them to fixate on new locations continuously. More importantly in the framework of this chapter, humans have intentions and goals to link the past with present with the aim of anticipating the future; animals interpret perceptual input by using their knowledge of images, sounds, actions, and objects, along with the perceptual operators that extract information from signals. Human actions, in particular, are goal driven, and they are guided not only by motor expectations but also by perceptual expectations (Rao and Ballard 1997; Sandini et al. 1993). Cognitive vision is an expectation-driven process, and in this sense, vision supports both the recovery of perceptual information to guide actions and the process of predicting the perceptual consequences of goal-directed actions. Currently, this debate about the nature of the perceptual process is no longer a philosophical nature as it has acquired

practical significance. As the field of cognitive robotics is evolving, practitioners and theorists are faced with basic questions: How should the visual system of the cognitive robot be structured? Specifically, should it be a black box delivering a 3D model of the scene in view along with labels for the objects and actions happening in the scene, or should it be structured differently, more in line with what biological systems do? If it ought to be the latter, precisely what would that be?

14.2 Is Perception (Only) an Inference Process?

Signal analysis is not enough to produce an understanding of a scene; rather, there must be some additional source of information beyond the images that can be used in the process of perception. The physicist von Helmholtz proposed that the additional knowledge is brought in through a process of inference—as we look at the world, we are also thinking about it. Since we are not aware of this thinking, he labeled it unconscious inference. Indeed, adding any form of knowledge to the signal processing can be considered reasoning or making an inference. The prior knowledge people bestow upon a scene is about the likely configurations of objects, events, and their relations, along with basic physics. Thus, perception interacts continuously with cognition at different levels of abstraction: it guides attention, constrains the search space for recognition, reasons over what is being perceived, and makes predictions about what will be perceived. This is an interactive bottom-up and top-down process; as visual (perceptual) information is anticipated from past experience and searched for through purposive actions, meanings emerge. This is what we call cognitive vision, which can be succinctly defined as a system implemented as a continuous exchange of information between perception and reasoning. It is a form of predictive vision in the sense that it does not simply rely on actions to optimize information acquisition; rather, actions are driven by perceptual expectations (How should I act to see my hand close to the object to grasp vs. how should I act to reach the object?).

14.3 Cognitive Vision: The Vision of the Middle Layer (in an Embodied Framework)

Low-level perception is traditionally thought of as feeding into a high-level knowledge database (KB) where inference can work (Marr 1982). However, in cognitive systems, perceptual outcomes can be predicted, and through embodiment they can be actively searched for through goal-directed actions. The goal is to obtain the expected sensory input; as such, actions are a means to achieving this goal. This generates the mutually supportive roles of perception and inference through actions and expectation. This occurs not simply at low and high levels but also at an intermediate or midlevel where input-driven, bottom-up signals are combined with top-down, expectation-driven signals. At this level, reasoning and perception talk with each other about objects, actions, events, and alternative possibilities in a kind of *internal cognitive dialogue* that loops between *prediction* (what the system expects perceptually) and *exploration* (how the system acts to verify if these predictions are being met). In this framework, “active inference” (Friston, Daunizeau, and Kiebel 2009) is the process of inferring which actions minimize the error

between the expected sensation and the resultant outcome. Thus, cognitive vision is the set of processes that processes real-time information and provides the perceptual hypothesis required to carry out this dialogue (the “predictive coding” stream of research by Rao and Ballard [1997] and the “learning to predict the next sensation” proposed by Tani and Nolfi [1999] and recently reported by Nagai [2019]).

With respect to computer vision, the peculiarity of cognitive vision is the extension of the concept of “processing visual data” beyond the concepts of “extracting visual features for real-time control” (as in reactive systems) to address how to “generate expected visual features supporting anticipatory behavior.”

Here are five rough *interaction paths* through which Vision (V) and Reasoning (R) can engage during an internal cognitive dialogue.

1. $V \rightarrow R$. This is the traditional perspective of first applying computer vision and then transferring the results to AI for reasoning.
2. $R \rightarrow V$. For example, “search for the scissors” starts with the concept “scissors” and invokes a visual search. When we perform a task and we follow a procedure, we continuously invoke this path.
3. $V \rightarrow R \rightarrow V$. For example, the vision system concludes that the activity taking place is “someone is cutting the tomato with a spoon,” and it communicates this to the Reasoner. Subsequently, the Reasoner finds it implausible and asks V to check again to determine whether the tool is a spoon or a knife.
4. $R \rightarrow V \rightarrow R$. For example, the Reasoner needs to know the number of cars in some location, and it initiates a counting search for the Vision system.
5. $R \rightarrow VV \dots V$. This amounts to imagining and envisioning a situation, action, or event.

As such, the interactions between V and R are many and complex, and it is not clear how one should best develop them.

The interactions between Vision and Reasoning can happen at different stages. First, it can happen at “later stages,” meaning that vision is running recognition procedures and producing symbolic information that it gives to the Reasoning process. Second, it can happen at “earlier stages,” meaning that Reasoning helps Vision by resolving the ill-posedness of visual processes. In the latter case, instead of vision performing only a bottom-up segmentation and recognition, additional knowledge can be introduced. For example, it is easier to segment an object with known attributes, such as delineating a “long red object” from a background, as opposed to generically segmenting the scene into surfaces using only bottom-up vision. Moreover, the interaction can also happen at the middle stages where, for example, knowledge about the action produces expectations for both objects and movements.

As discussed in Vernon (2006), cognitive vision can be defined in terms of its generic functionalities (i.e., detection, localization, recognition, categorization, and understanding of an object or event), its nonfunctional attributes (i.e., purposive behavior, adaptability, and anticipation), and how it supports the acquisition, storage, and use of knowledge (i.e., learning, memory, and deliberation). This final element is considered “the key differentiating characteristic of Cognitive Vision vis-à-vis Computer Vision” (Vernon 2006).

How should cognitive vision be implemented? Unfortunately, this question does not have a clear answer. It would be useful if existing cognitive architectures could be used to implement

the exchange between vision and reasoning. However, existing cognitive architectures do not consider this exchange at the core (Laird 2012; J. Anderson 2007; Kotseruba and Tsotsos 2020; Vernon, von Hofsten, and Fadiga 2010). Thus, instead of devising an architecture for cognitive vision, it would make sense to list a number of principles or attributes of cognitive vision that should be considered in the development of any architecture. Fundamentally, the architecture has to be a “message-passing architecture.”

14.4 Principles of Cognitive Vision

14.4.1 Cognitive Vision Is to Support Prospection

A peculiarity of cognitive vision is that it is inherently linked to the interactions of the perceptual agent with the environment. This includes other agents and, within the scope of this book, it specifically emphasizes the interactions between the embodied agents that can not only analyze the environment but also change the environment through their actions. The distinction between analyzing and changing brings about the proposed organization of the human visual system into two information-processing streams that, originating in the occipital cortex, project dorsally to the parietal cortex and specialize in spatial vision and ventrally to the temporal cortex, specializing in object recognition. The original evidence of the so-called what-and-where organization was proposed initially in Ungerleider and Mishkin (1982) based on behavioral studies in monkeys. In one study, lesions of the *what* (ventral) pathways produced an inability to discriminate objects while leaving intact visuospatial tasks such as visually guided reaching. In turn, lesions in the *where* (dorsal) pathway did not affect visual discrimination but severely affected performance in visuospatial tasks. In their original paper, Ungerleider and Mishkin (1982) proposed that the ventral processing stream mediated the visual recognition of objects (“what” an object is), while the dorsal processing stream mediated the appreciation of the spatial relationships among objects and the visual guidance toward them (“where” an object is). Goodale and Milner (1992) proposed a modification to this model, emphasizing “perception” versus “action” for the ventral and dorsal processing streams, respectively. However, this separation cannot be interpreted as a chiasm between the areas involved in the execution of actions and the areas devoted to recognizing objects because, in general, the view of the brain as a collection of areas connected statically has evolved in the vision of the brain as a dynamical system in which individual regions are functionally diverse and used and reused in many different tasks across the cognitive domain (M. Anderson 2014). Certainly, spatial information is not segregated into the dorsal pathway, but it is closely integrated with “object processing” both in terms of relational dimensions as well as the position of objects in the environment. The current view is that “the ventral and dorsal pathways treat objects and space differently, but they cannot treat them separately” (Connor and Knierim 2017). From the point of view of prospection, it is nonetheless true that expectations play an important role in the object’s localization and recognition (as demonstrated by Biederman, Mezzanotte, and Rabinowitz [1982]; see figure 14.1). Actions play a fundamental role in building knowledge about the object’s properties; in turn, these properties are retrieved from memory to both act on and recognize objects. The cognitive processes exploiting this shared memory do not distinguish between anticipating the occurrence of



Figure 14.1

Example of position violation supporting the view that anticipation is involved in object recognition. *Source:* From Biederman, Mezzanotte, and Rabinowitz 1982.

objects and events. Therefore, even though this chapter is focused upon actions and behavior and less upon object categorization and recognition, we believe that the cognitive principles that we describe apply to the information carried through the ventral as well as the dorsal visual pathways.

From the cognitive vision perspective, embodied interaction adds the need to use vision “to control actions” as well as the use of vision “to anticipate actions.” Cognitive vision supports prospection both through detection and recognition. Within the “cognition” framework, prospection refers to the ability to anticipate the outcome of the (inter)actions, including its sensory outcome. The visual system must be able to “imagine” (anticipate, synthesize) what it is going to see as a result of a given action and, during actions, to “monitor” whether the sensory expectations are being met (predictive vision; Sandini et al. 1993). It is worth noting here that the need for this “anticipatory” role of vision has been proposed by neuroscientists such as Alain Berthoz (1997) and Marc Jeannerod (2001) in his “mental simulation theory,” as well as by computational neuroscientists such as Rao and Ballard (1997). The behavioral difference here is that the agent can plan its actions so they cause a specific sensory outcome. If what I see is an object, I should see a different “optical flow” pattern when I push it or grasp it, and I plan the action to verify the matching of the “imagined” (desired) pattern with the one being generated contingently (Sandini et al. 1993). That is to say, through purposively planned actions and monitoring of anticipated sensory outcomes, behavior-related visual features are used to segment objects in a visual environment, as shown in figure 14.2. Objects are not identified only by edges and blobs.

This mapping from a reactive *action-perception* loop to an anticipatory *exploration-prediction* loop (in which actions are planned and sensory outcome is anticipated) does




Figure 14.2

Object segmentation through purposive action. During a pushing action (a), motion information (optical flow pattern, b) is extracted to detect the instant of contact and to segment the approaching hand (c) and the object pushed (d). *Source:* Adapted from Sandini et al. 1993.

not diminish the importance of extracting and measuring other “visual features” per se; rather, it extends their use to a generative model. An important consequence of this cause-effect link between perception and action is the different ways they represent shape. In the anticipatory exploration-prediction loop, the emphasis shifts from the features used to describe geometric features (such as generalized cylinders) to action-based, proactive features. A striking example of the latter are the “canonical neurons” that code the shape of an object as a function of the actions best suited to grasping it (Fadiga et al. 2000). In this way, shape is defined by “grasp type,” and conversely, the hand is preshaped during reaching actions to encode and anticipate the shape of the object that will be grasped (Campanella, Sandini, and Morrone 2011; Gori et al. 2011). The relevance of these and other anticipatory features in relation to human-human and human-robot interactions is the subject of the next section.

14.4.2 Cognitive Vision Is to Support Human-Robot Interaction

A vast amount of literature discusses how anticipation affects vision in both the recognition and categorization of objects and events as well as visually guided behavior. However, considering the primary role of cognition on social behavior, here we focus on the special case of the ability to anticipate the goals and intentions of a partner agent during social interactions. In this framework (referred to as “social cognition”), vision is an important channel (before physical contact occurs vision is the only channel) used to synchronize the activities during the execution of collaborative tasks and to derive the kind of shared goals and intentions necessary for two agents to work *together* instead of simply working *next to each other*. In this case the peculiarities of cognitive vision are not limited to the use of visual information to “control” movements (e.g., the trajectory of the hand toward the object to be grasped) or the *forces exchanged* (such as when multiple entities jointly handle a box). Rather, they refer to the use of vision as a communication channel through which the interacting agents are exchanging “messages” (Sciutti et al. 2012; Sandini, Sciutti, and Rea 2019). Therefore, the role of vision extends from anticipating the outcome of the agent’s own actions (as described in section 14.4.1) to anticipating (understanding) the outcome of the partner’s actions, including their intentions and goals. The special case of an anthropomorphic partner is particularly interesting because it offers the possibility of using humanoid robots to study aspects of social cognition in humans that are still unknown and that cannot be investigated in other ways (Sandini 1997; Sciutti et al. 2015). Here, we will focus on the special case of an anthropomorphic agent even though what we are reporting may be applied to systems with different degrees of anthropomorphism and not only to humanoid robots. In this case, behavior-based communication originates primarily from bodily features such as the physics and topology of biological sensors and actuators.

In relation to the kinds of messages that can be exchanged visually between embodied agents, we first need to make an important distinction between implicit and explicit messages. Implicit (covert) social signs are expressions of a physical property of the body or of a general law of physics (e.g., gravity), and they are not under voluntary control. In contrast, explicit (overt) social signs are voluntarily controlled, and they include gesticulation language and culture-related movements (Sciutti et al. 2018). In both cases many segments of the human/robot body can be involved in producing such social signs: eyes, hands, heads, face, appearance, whole-body posture, and so on (messages are limited to those that can be exchanged visually). In relation to using cultural-based, overt gestures, a vast amount of computer-vision literature exists in relation to gesture recognition in general and hand signs in particular (Wu and Huang 1999; Al-Shamayleh et al. 2018), even if we disregard all the methods used for movement studies based on wearable accelerometers, reflective markers, and active sensors.

In this article we concentrate on the role vision plays in anticipating one’s partner’s goals and intentions on the basis of implicit, body-based messages specific to human social interaction (and, to some extent, also to other social species). By “body-based,” we are here referring to the messages embedded in some fundamental, physical property of the human body that are inherited as a result of evolution.

The first, most obvious message is the direction of gaze, which makes an actor’s intention explicit by displaying her “region of interest” (see figure 14.3). This is an anticipatory



Figure 14.3

Without the subject being aware, the iCub's behavior was driven only by the direction of the subject's gaze making explicit her intention to reach for the left or right hand during an interaction task. *Source:* From Palinko et al. 2016.

feature, as gaze direction anticipates grasping actions and is used by the actor to guide reaching (Gandolfo, Sandini, and Bizzi 1996; Sandini, Metta, and Konczak 1997; Flanders, Daghestani, and Berthoz 1999); moreover, it is used by the observer to anticipate the intentions of the actor (Johansson et al. 2001). This holds true even in young infants (Falck-Ytter, Gredebäck, and von Hofsten 2006). See Gredebäck and Falck-Ytter (2015) for a recent review on the effect of eye movements during action observation.

In the case of mutual gaze, fixation has an even more important social effect in establishing a preferential communication channel between two agents. Even if one's gaze can be actively controlled and one's attention is not necessarily linked to the direction of the gaze, eye movements are ultimately needed because of the space-variant nature of our retina, and as such, they are a fundamental implicit visual signature of the agent's internal state. The role of the gaze is well known in humans and great apes (Tomasello et al. 2007; George and Conty 2008), and it has been exploited in different areas of research with specific emphasis on joint attention (Doniec, Sun, and Scassellati 2006), eye contact (Palinko, Sciutti, et al. 2015; see figure 14.4), human-robot engagement (Hall et al. 2014), and turn-taking.

In addition to gaze direction, the way humans (as well as other biological systems) move has regularities that are ultimately linked to the very nature of their muscles (a true embodiment constraint). If these regularities can be perceived visually (i.e., if they are within the range of the signal that vision can "see"), they may be related to specific visual features. As stated very clearly by Jeannerod and Jacob (2005), "Not only is what one can do shaped by what one perceives, but also conversely what one can do shapes what one can perceive." It is on the basis of these regularities that humans (and other biological systems; Regolin, Tommasi, and Vallortigara 2000) can distinguish biological motion from the movements generated by other sources (e.g., a tree's branches being moved by the wind, vehicles traveling, balls rolling, and so on). The ability to detect biological motion, which is present from birth, is one of the powerful stimuli that allow humans to develop social cognition.




Figure 14.4

Contingent detection of eye contact by the iCub was used to time turn taking during a dictation task. The performance improved with respect to timing related to word length. The subject was unaware of the difference but reported a more natural condition in the “contingent” situation. *Source:* Adapted from Palinko, Rea, et al. 2015.

For example, a signature of biological motion known as the two-thirds power law is applied to movements of a body segment (e.g., the hand during drawing or writing or the knee or ankle during walking); this reflects the relationship between angular velocity and the curvature of the trajectory (Lacquaniti, Terzuolo, and Viviani 1983; Viviani and Terzuolo 1982; see also Richardson and Flash [2002] for a comparison with other descriptions and generalizations). This parameter can be visually measured and used to identify and segment a moving biological body independently of its shape and color (Vignolo et al. 2017). Thus, it is independent of skin color or of the occlusions hiding one’s body (a shadow of a body segment can still be detected in the same way). Of course, this is just the initial segmentation phase of understanding gestures and movement; nonetheless, it is important for its unique biological signature (see figure 14.5).

Another signature of biological motion specific to ballistic movements (such as reaching) is the velocity profile of the end effector, which has a particular bell-shaped profile (Abend, Bizzi, and Morasso 1982; Morasso and Mussa Ivaldi 1982). It is worth mentioning that both these low-level visual measures are not only useful for discriminating biological systems; they can also be further exploited to anticipate the intentions and goals of the partner and the properties of the manipulated objects that are, apparently, not accessible through vision (e.g., anticipating the instant in time when the hand will be reaching a target point in space and where that target is located in the image). An interesting example of this is the human ability to estimate the weight of an object being lifted by another person.




Figure 14.5

Segmentation based on biological motion obtained through optical flow computation and identification of blobs moving according to the two-thirds power law constraint. *Source:* Adapted from Vignolo et al. 2017.

This ability, which is present in adults (Hamilton, Wolpert, and Frith 2004; Hamilton et al. 2007; Senot et al. 2011), develops during the first six to ten years of one's life (Sciutti, Patanè, and Sandini 2019) in synchrony with children's ability to exploit weight to successfully execute manipulation tasks (Wang, Williamson, and Meltzoff 2018).

The motion parameters that elicit the visual signature used by the observer to estimate the object weight were identified by modulating the kinematic parameters of a robot actor lifting objects with different weights and comparing the results with the estimated weights elicited by a human actor. The results showed that the relevant features are related to velocity profile (Sciutti et al. 2014; see figure 14.6). In this way we can literally see the forces exerted by a partner even without (or before) having any physical contact.

Another important aspect of the visual features described here that have a fundamental role in interaction is that they can be used to anticipate the outcome of one's own action (e.g., the instant in time I will touch the object), and they can be easily extended to anticipate the effects of the actions of other conspecifics (meaning the agents with similar structures and/or similar motion parameters). In fact, both measures (i.e., the two-thirds power law and the velocity profile) are independent of the relative positions of the observer and the actor (thus, they are independent of perspective transformation). As a special case, the actions described on the basis of these parameters are invariant to mirror transformation, and they could very well be the elementary visual measures that contribute to the activation of mirror neurons. In this way, the expected velocity profile (and the related behavioral measures) of one's own hand or someone else's hand reaching for an object are the same, and they can be used to anticipate both the position of the targeted object and the expected time to contact.

Of course, this does not solve the issue of action understanding, which requires a hierarchical representation of motor symbols to be handled, memorized, and recalled. However, it offers an alphabet (or part of an alphabet) with important invariances that can be exploited. For example, one can exploit what is learned by an agent by looking at one's own actions, not only to learn to move but also to learn to understand.

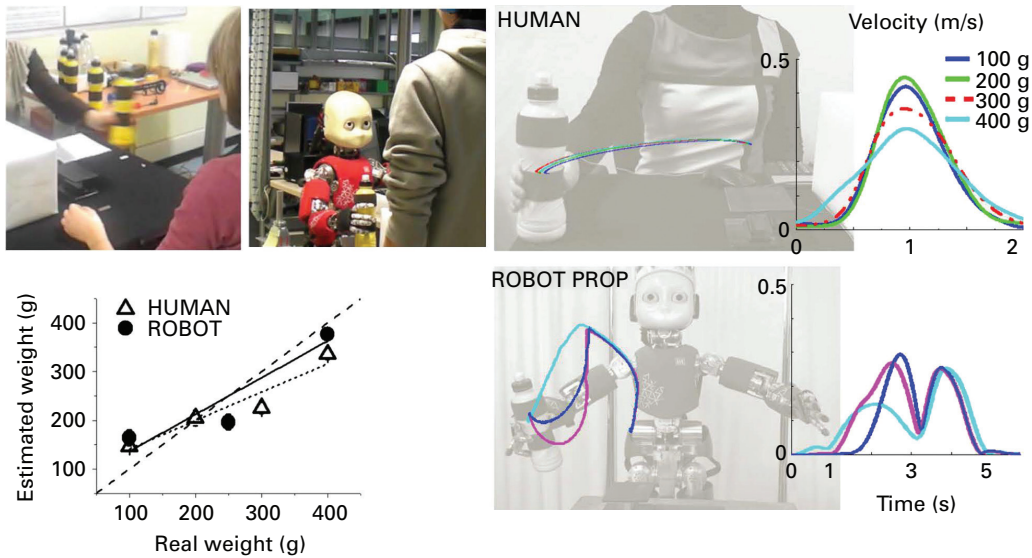


Figure 14.6

In this experiment, subjects were asked to report the weight of objects after observing both an actor and a robot lifting, transporting, and putting in place a set of bottles that were visually identical but of different weights. The possibility of fine-tuning and modulating the motion parameters of the robot (using the robot as a stimulus) and comparing the outcome with human observation allowed researchers to identify the vertical velocity of the action and its duration as the most relevant kinematic parameters for weight estimation on the basis of action observation. See Sciutti et al. (2014) for more details.

Another aspect of action understanding (which goes more directly into understanding one's state of mind) that can be supported by visual measures is the *style of action*, which conveys indirectly the internal state of the actor (e.g., actions performed while calm vs. the same action performed while angry or impatient; Di Cesare et al. 2019; Vannucci et al. 2018).

14.4.3 Cognitive Vision Involves Language as an Attention Mechanism

One can connect perception with reasoning through knowledge-based engineering and through the use of language—basically, using language to reason. The field of computer vision has led to interest in introducing additional higher-level knowledge about image relationships into the interpretation process (Farhadi et al. 2010; Forsyth et al. 2009; Gupta, Kembhavi, and Davis 2009). Although existing work acquires this additional information from captions or related texts, one could use more advanced techniques to obtain additional high-level information.

Computational linguists have an interest in lexical semantics—that is, conceptual meanings of lexical items and how these lexical items relate to each other (Cruse 1986). They have also created resources through which we can obtain information about different concepts, such as cause-effect, performs-functions, used-for, and motivated-by. For example, the WordNet database relates words through synonymy (words having the same meaning, like “argue” and “contend”) and hypernymy (“is-a” relationships, as between “car” and “vehicle”), among many other relations (Miller and Fellbaum 2007). Linguistics has also created large text corpora and statistical tools that enable us to obtain probability distributions

for the co-occurrence of any two words, such as how likely a certain noun co-occurs with a certain verb.

Using these linguistic tools, we can aid vision to build better systems for interpreting images. One way is to use linguistic information as a contextual system that provides additional information to interpretations; this is already utilized in some multimedia systems. Certain objects are likely to co-occur; for example, “tables” often co-occurs with “cups” and “spoons.” But if we consider vision an active process, there is more than just observation. Let’s say you are in a kitchen. Because you have prior knowledge about kitchens, their structure, and the actions that take place in them, and because a large part of this knowledge is expressed in language, we can utilize this information during our visual inspections. A knife in the kitchen will most probably be used for cutting a food item, so one’s vision can look for it. In this way, language acts as a high-level prior knowledge system that aids perception. Moreover, let’s say you observe someone picking up the knife and putting it in a drawer. You know that the object is not gone; rather, it is just hidden from sight. In this case, language acts as a part of the reasoning process. When humans interpret a visual scene, we fixate upon some location and recognize the nouns, verbs, adjectives, adverbs, and prepositions that are part of that location. Because our linguistic system is highly structured, our recognition produces a large number of inferences about what could be happening in the scene. Subsequently, when we fixate upon a new location, the same process is repeated. In other words, language acts as part of the attention mechanism of humans.

14.4.4 Cognitive Vision Is Supported by a Question/Answer Mechanism: The Cognitive Dialogue

During the process of vision, our visual procedures interact with language processes and motor actions. Low-level perceptual object features and/or movements, high-level knowledge, and our overall goals guide our attention. The processes that recognize objects and actions interact continuously with prior knowledge that enables people to formulate expectations; in turn, this constrains the recognition search space. Reasoning is used to analyze visual input and, if needed, to correct the visual recognition toward the solutions that make sense within the specific context. This dynamic interaction of cognitive processes is generally agreed upon, but it has not yet been implemented computationally.

We suggest that this interaction should be implemented as a dialogue computationally so that intelligent systems can achieve scalable visual scene analysis. For example, let’s say the goal is to produce a semantic description of the scene in view. This can be achieved by having the Reasoner (R) and the visual processes (V) engage in some form of a cognitive dialogue through language. The R can ask the V a number of questions, such as: Is there a *noun* in the scene? Where is it? What is next to the noun? Where did the agent that performed *actionX* go afterward? By allowing the R to ask questions and to receive answers, and then by repeating this process, we bring forward the whole power of language in the semantic analysis. This is something that has not been possible before. If we also include the motor processes and the auditory processes (AP), the cognitive dialogue integrates perception, action, and cognition.

14.4.5 Cognitive Vision Uses Linguistic and Optimization Tools

Consider the problem of activity description, as in this sentence: A man cuts the tomato with a knife on the counter (i.e., assume that the cognitive system is observing a man cutting a tomato with a knife, and it needs to come up with the linguistic description “a man cuts a tomato with a knife on the counter”). Language tools should provide us with two kinds of information. First, we need information about the possible quantities within a certain context. In the example above, assuming we know that we are processing a kitchen scene, language should provide the possible objects and verbs that are generally used in this setting. Current computer vision applications deal with predefined data sets. However, knowledge databases can provide this new information. The Praxicon Project (Poeticon Project 2012; Pastra et al. 2011) is a resource that contains knowledge of commonsense everyday activities. This lexical database, obtained by reengineering WordNet, provides pragmatic relations about verbs and objects and also provides algorithms that we can use to query domain-specific knowledge (for example, if we want to obtain the quantities involved in cooking a particular meal). Second, we need language tools to provide us with the contextual relations of different quantities, such as that “knives” are possibly used for “cutting,” and this activity is often performed “on tables.” Classical linguistics can build domain knowledge of this kind, and it can provide information on whether or not certain combinations are plausible. Statistical language tools that access large text corpora can provide statistics on the co-occurrence of the different quantities in certain domains. Subsequently, we can use this statistical language information, along with the statistical information gained from the visual recognition, with classifiers to optimize for scene interpretation.

For an example, see Teo et al. (2012) for the use of the statistical language approach for interpreting actions in video, where the probabilities of the co-occurrence of quantities were obtained from a large text corpus to generate a sentence description. In addition, the lexical database approach has been demonstrated to enable a robot to observe humans performing actions and then subsequently to create descriptions that will allow the robot to execute similar actions (Summers-Stay et al. 2012; Yang et al. 2014). Interesting questions arise when we realize the dialogue for active agents and then construct the models dynamically. As the dialogue continues the construction, the additional knowledge introduced into the language space changes the expectation for other concepts. Similarly, knowledge creates expectations in the visual space, and thus, it constrains the search space for object and action recognition. For example, if there is a high probability during the dialogue that the human is performing a cutting action, the vision module will not need to run every object classifier to identify the tool used for cutting; rather, it will examine a small set of cutting tools to determine which tool is being used. Going even further, instead of applying object classifiers it could instead apply procedures that check for the appearance of cutting tools.

The cognitive dialogue is also well suited for action interpretation because actions are compositional in nature. Starting from the simple actions that occur on a part of the body, we can compose actions from several limbs to create more complex actions; further, we can combine a sequence of simple actions into activities (figure 14.7). Language can be used to further enhance action recognition at the higher levels using the compositions from lower levels. Moreover, language can be used to enforce temporal and logical constraints

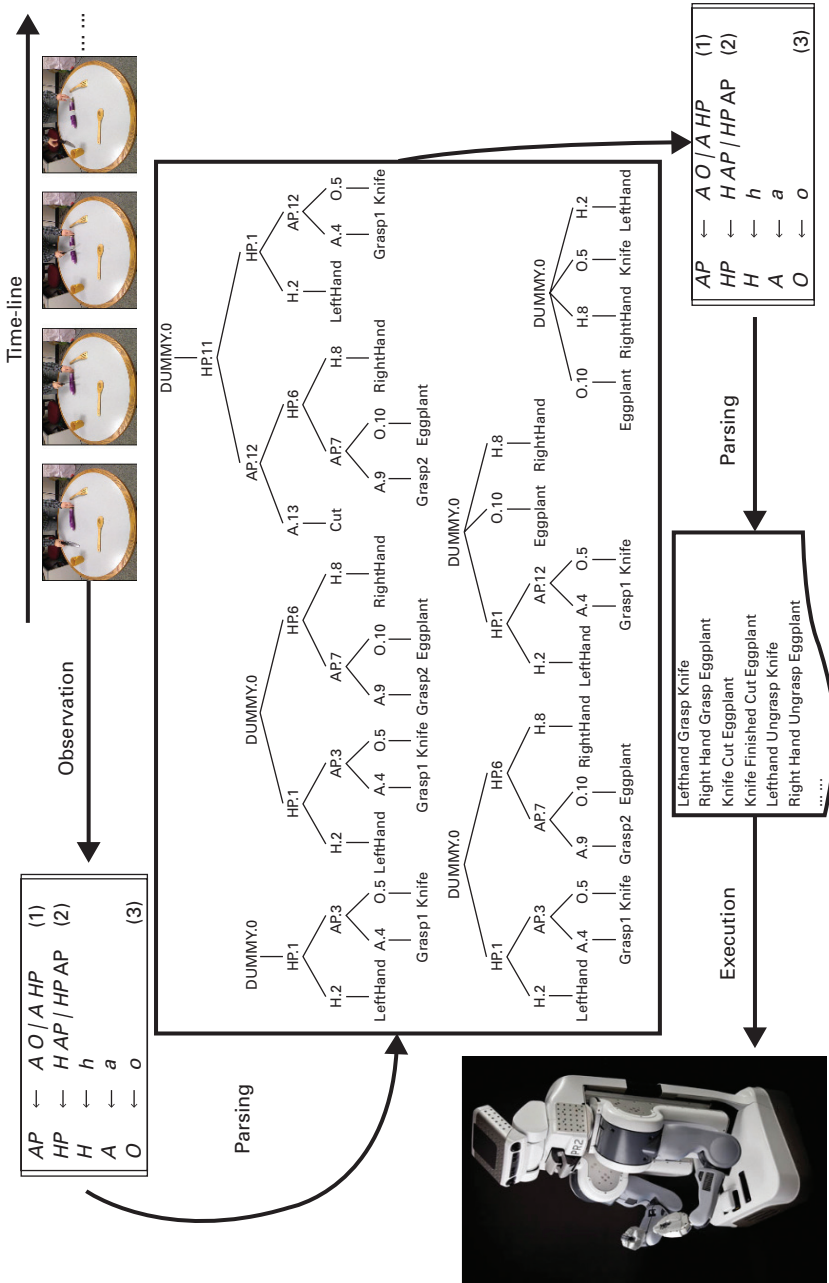


Figure 14.7

Using an action grammar, a video of a person cutting an eggplant on the table (*top right*) is turned into a sequence of parse trees (*middle*). Using the parse trees, the activity can be turned into a sequence of commands that a robot can execute. *Source:* From Yang et al. 2014.

on how actions can be chained together by using a grammar of action (Pastra and Aloimonos 2012) that binds sensorimotor symbols (hands, arms, body parts, tools, objects) with language. In this case the Reasoner R will work across all levels, from the bigrams of actions to inferring the most likely activity that could be occurring by examining the sequence of such bi-grams, using large corpora.

14.4.6 Cognitive Vision Uses Knowledge-Based Control

An essential component of the vision knowledge/language dialogue is that it should guide the attention to expected objects, their locations, and their attributes and to the actions in the scene. Thus, we need models for the attention mechanism that will predict the order of fixations, specifically to what and where we should allocate the computational resources. As an example, Yu et al. (2011) proposed a way to control cognitive dialogue using information theory. At all times the system has a goal. This goal can be as simple as recognizing a scene by the objects in it or as complicated as recognizing an activity that is described by many quantities. At each time t , the system must utilize what it already knows in order to determine the optimal question to ask at step $t + 1$.

The researchers formulated this technique using Bayesian estimation. The probabilities involved come from the accuracy of the visual detectors; the importance of the individual quantities in describing the scene or activity is derived from language. At step $t + 1$, the criterion for selecting the appropriate quantity is to maximize the information gained about the scene/activity recognition as a result of the response of the added quantity detector. This can be modeled by considering the KL divergence between the probability distributions of detecting the activity on the basis of the quantity detectors at step t and the probability distribution of detecting the activity upon the addition of a new quantity detector. By adding criteria for how to start the process (for example, such as always attending first to moving humans) as well as criteria for finishing the process, we obtain a systematic way of carrying on the dialogue.

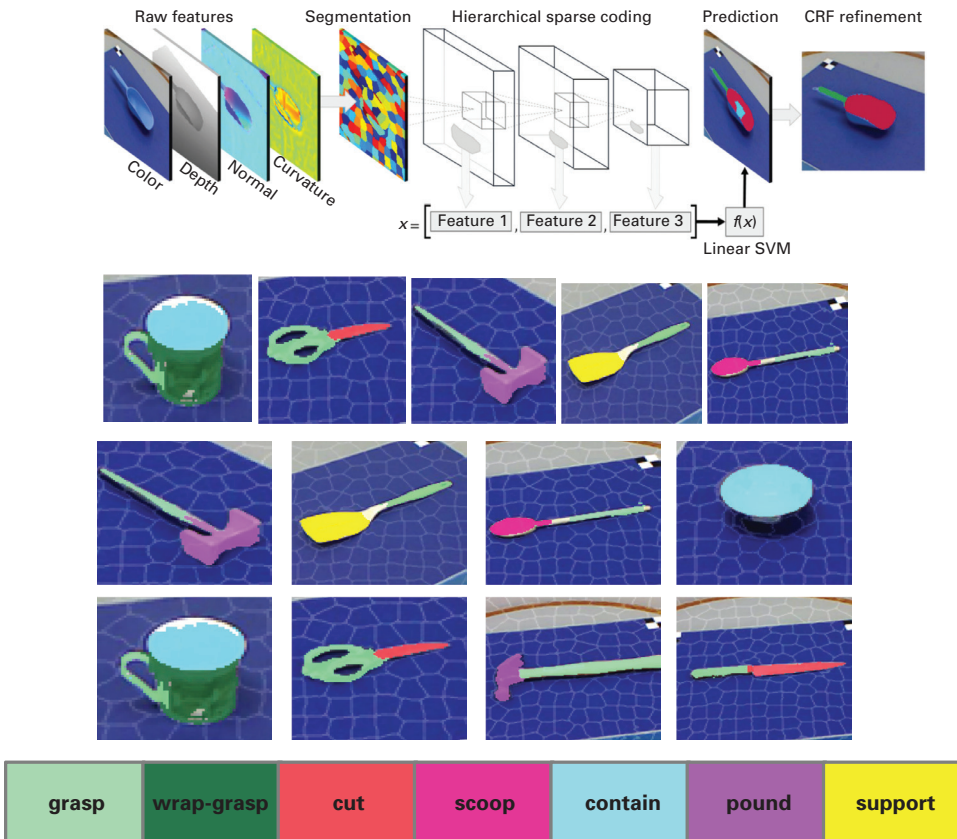
14.4.7 Cognitive Vision Senses “Parts of Speech” Operators

Because the dialogue is carried out in language, it is important for vision to find and recognize nouns (objects), verbs (actions), adjectives (attributes), adverbs (manner of the action), and prepositions (spatiotemporal relations). In that way, V can search for a “red object,” for a “long blue object,” for “the object to the left of the noun,” or for “the action before the tomato was picked up.” In this way the operators are the grounding of language in perception. For the past several years, the community at large has been working on various aspects of this problem, such as object recognition, attribute-based recognition, action recognition, and prediction (Fermüller et al. 2018).

The cognitive dialogue offers a new way of recognition by making it similar to the twenty-questions game; we can ask questions about the object related to its affordances or properties (Jamone et al. 2016; Myers et al. 2015; see figure 14.8) in order to recognize the object. Segmentation becomes an important operation since it is needed to infer adjectives related to shape (Mishra and Aloimonos 2009; Mishra et al. 2012; see figure 14.9 and figure 14.10). However, the recognition of verbs remains a challenge. There has been progress with a number of data sets; however, we are still far from achieving robust action recognition procedures based on visual information alone (Carreira and Zisserman 2017;

Affordance of Object Parts from Geometric Features

Using hierarchical matching pursuit (M-HMP)

**Figure 14.8**

Using machine-learning techniques, we can employ supervised learning to “color” each pixel with its associated affordance. *Source:* Myers et al. 2015. *Bottom:* Results from the application of the learned classifier to the images of tools.

Damen et al. 2018; Gu et al. 2018; Caba Heilbron et al. 2015; Soomro, Zamir, and Shah 2012; Sigurdsson et al. 2016).

14.4.8 Cognitive Vision Is Supported by Deep-Learning Techniques

Cognitive vision amounts to a set of processes; some of these are geometric or photometric, while others interact with reasoning, planning, language, and the motor system. Many of these processes could be implemented in a deep-learning framework. For example, take the process of segmentation described in the previous section, in which the analysis is happening in a (log) polar space. In this case we can develop deep networks that are able to learn the transformation from Cartesian to polar space (Esteves et al. 2018). In a similar fashion, using supervised convolutional neural networks (CNNs), we can learn the appearance of different hand grasps that can aid in recognizing manipulation activities (Yang

Compositionality and Depth Boundaries

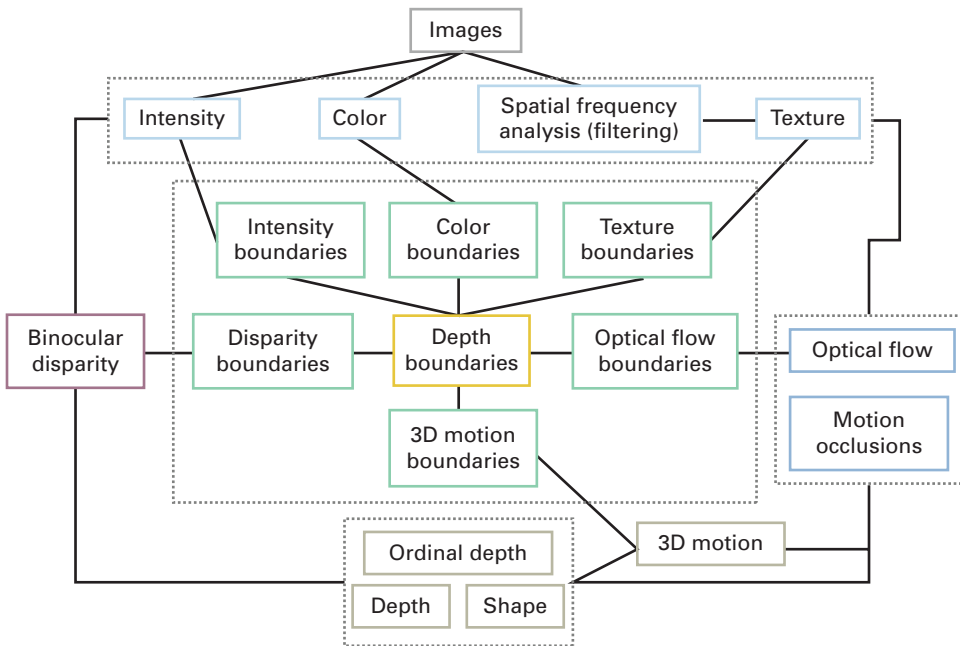


Figure 14.9

Cognitive vision focuses on the segmentation problem. All low-level cues participate in finding the boundaries in a compositional manner. *Source:* From Ogale and Aloimonos 2007.

et al. 2016). Overall, many of the geometric aspects of cognitive vision can be formulated into a deep-learning framework.


Along a different path, recent approaches by researchers have utilized the power of deep learning to investigate attention, prediction, and semantics. See, for example, Vondrick, Pirsivash, and Torralba (2016) on anticipating visual representations from unlabeled video. Regarding semantics, the entities we recognize in a scene are structured in the so-called scene description graph (SDG; Aditya et al. 2018), which could be learned using deep learning. See, for example, the work of Xu et al. (2015) on neural image caption generation with visual attention, which takes one from vision to language.

14.5 Conclusion

It is important to stress that cognitive vision does not exist in isolation as a mechanistic system that learns to detect what is where; this is in direct contrast to how vision is predominantly studied today. Since the seminal works of Held and Hein (1963), Stein and Meredith (1993), and Milner and Goodale (1995) and the stream of experiments on the visual coding of space and actions by Graziano, Yap, and Gross (1994), Graziano et al. (2002), and Fadiga et al. (2000), it has become evident that, apart from the very early processing stages, cognition operates through a unified representation wherein vision and other sensory modalities are dynamically merged through action. In this sense, cognitive

Segmenting a fixated object

Step 1: Cartesian to polar with fixation as the pole



Step 2: Find the optimal cut through the polar map




Figure 14.10

Top left: The edges of the image shown at the bottom right with a green fixation point. By turning this into a polar map, segmentation amounts to finding a top-down contour. *Source:* Mishra et al. 2012.

vision has to be considered part of an intelligent system that reasons and acts. Thus, it can ask questions beyond what and where—such as why, how, who, and many other questions (Amy and Song-Chun 2013; Verschure 2012). Cognitive vision does not only address how to extract information from images to control actions; it also considers how to synthesize visual information by anticipating the effects of actions. As such, cognitive vision introduces more interesting questions, and it points to a new theory for the integration of intelligent systems with perception.

Additional Reading and Resources

- Cognitive vision can be seen in some recent papers of the current literature, under the heading of question answering or visual-question answering or visual search. Indeed, to answer a question related to an image (or a video) one would need procedures for search (matching words to parts of the scene (image or video), procedures for recognizing actions as well as their constituents (objects, tools, agents), and procedures for predicting activi-

ties. In addition, some form of commonsense background knowledge is required. These operations are examples of cognitive vision. See, for example: 1) Aditya, Somak, Yezhou Yang, and Chitta Baral. 2018. “Explicit Reasoning over End-to-End Neural Architectures for Visual Question Answering.” In *32nd AAAI Conference on Artificial Intelligence*, 629–637. Menlo Park, CA: AAAI Press; 2) Goyal, Y., T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. 2017. “Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering.” In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. New York: IEEE.

- This work presents the first effort to generate commonsense captions directly from videos in order to describe latent aspects such as intentions, attributes, and effects: Fang, Zhiyuan, Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. 2020. “Video-2commonsense: Generating Commonsense Descriptions to Enrich Video Captioning.” In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 840–860. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.61>.
- Answering questions related to images or video is a problem of cognitive vision, and research on the issue is expected to flourish in the near future. The relationship between vision, action, and language will be central to this enterprise: Moens, M.-F., K. Pastra, K. Saenko, and T. Tuytelaars. 2018. *Vision and Language Integration Meets Multimedia Fusion*. New York: IEEE.
- This book chapter by David Vernon is a nice starting point introducing “cognitive vision” in the context of visual perception and explaining its role in the context of cognition: Vernon, David. 2006. “The Space of Cognitive Vision.” In *Cognitive Vision Systems*, edited by H. I. Christensen and H. H. Nagel, 7–24. Berlin: Springer-Verlag.
- The relationship between cognitive vision and embodiment is the central theme developed here: Vernon, David. 2008. “Cognitive Vision: The Case for Embodied Perception.” *Image and Vision Computing* 26:127–140.
- This chapter by Sandini et al. is a good starting point for a review of works on the use of vision to understand communication signs exchanged through gestures: Sandini, G., A. Sciutti, and F. Rea. 2019. “Movement-Based Communication for Humanoid-Human Interaction.” In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 2169–2197. Dordrecht: Springer Netherlands.
- This Introduction to Cognitive Robotics course (www.cognitiverobotics.net) has several lectures (from fourteen to twenty) devoted to robot vision, with a comprehensive description of the important aspects of vision, which are then put into the context of cognition and cognitive architectures.
- The most commonly used software library for vision research is OpenCV: opencv.org.

References

- Abend, William, Emilio Bizzi, and Pietro Morasso. 1982. “Human Arm Trajectory Formation.” *Brain: A Journal of Neurology* 105 (2): 331–348.
- Aditya, Somak, Yezhou Yang, Chitta Baral, Yiannis Aloimonos, and Cornelia Fermüller. 2018. “Image Understanding Using Vision and Reasoning through Scene Description Graph.” *Computer Vision and Image Understanding* 173:33–45. <https://doi.org/10.1016/j.cviu.2017.12.004>.

- Al-Shamayleh, Ahmad Sami, Rodina Ahmad, Mohammad A. M. Abushariah, Khubaib Amjad Alam, and Nazeen Jomhari. 2018. "A Systematic Literature Review on Vision Based Gesture Recognition Techniques." *Multimedia Tools and Applications* 77 (21): 28121–28184. doi:10.1007/s11042-018-5971-z.
- Amy, Fire, and Shu Song-Chun. 2013. "Using Causal Induction in Humans to Learn and Infer Causality from Video." *Proceedings of the Annual Meeting of the Cognitive Science Society* 35. <https://escholarship.org/uc/item/4ng247kx>.
- Anderson, John R. 2007. *How Can the Human Mind Occur in the Physical Universe?* Oxford Series on Cognitive Models and Architectures. Oxford: Oxford University Press.
- Anderson, Michael L. 2014. *After Phenology: Neural Reuse and the Interactive Brain*. Cambridge, MA: MIT Press.
- Berthoz, Alain. 1997. *Le Sens Du Mouvement*. Paris: Editions O. Jacob.
- Biederman, Irving, Robert J. Mezzanotte, and Jan C. Rabinowitz. 1982. "Scene Perception: Detecting and Judging Objects Undergoing Relational Violations." *Cognitive Psychology* 14 (2): 143–177.
- Caba Heilbron, Fabian, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. 2015. "ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 961–970. New York: IEEE.
- Campanella, Francesco, Giulio Sandini, and Maria Concetta Morrone. 2011. "Visual Information Gleaned by Observing Grasping Movement in Allocentric and Egocentric Perspectives." *Proceedings of the Royal Society B: Biological Sciences* 278 (1715): 2142–2149.
- Carreira, Joao, and Andrew Zisserman. 2017. "Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308. New York: IEEE.
- Connor, Charles E., and James J. Knierim. 2017. "Integration of Objects and Space in Perception and Memory." *Nature Neuroscience* 20 (11):1493–1503. doi:10.1038/nn.4657.
- Cruse, D. Alan. 1986. *Lexical Semantics*. Cambridge Textbooks in Linguistics. Cambridge: Cambridge University Press.
- Cyc. 2014. <https://www.cyc.com>.
- Damen, Dima, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, et al. 2018. "Scaling Egocentric Vision: The Epic-Kitchens Dataset." In *Proceedings of the European Conference on Computer Vision*, 720–736. Munich: Springer.
- Di Cesare, G., C. Pinardi, C. Carapelli, F. Caruana, M. Marchi, M. Gerbella, and G. Rizzolatti. 2019. "Insula Connections with the Parieto-Frontal Circuit for Generating Arm Actions in Humans and Macaque Monkeys." *Cerebral Cortex* 29 (5): 2140–2147.
- Doniec, Marek W., Ganghua Sun, and Brian Scassellati. 2006. "Active Learning of Joint Attention." In *2006 6th IEEE-RAS International Conference on Humanoid Robots*, 34–39. New York: IEEE.
- Esteves, Carlos, Christine Allen-Blanchette, Xiaowei Zhou, and Kostas Daniilidis. 2018. "Polar Transformer Networks." International Conference on Learning Representations, ICLR 2018. <https://openreview.net/pdf?id=HktrIUlAZ>.
- Fadiga, Luciano, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzolatti. 2000. "Visuomotor Neurons: Ambiguity of the Discharge or 'Motor' Perception?" *International Journal of Psychophysiology* 35 (2–3): 165–177.
- Falck-Ytter, Terje, Gustaf Gredebäck, and Claes von Hofsten. 2006. "Infants Predict Other People's Action Goals." *Nature Neuroscience* 9 (7): 878–879.
- Farhadi, Ali, Mohsen Hejrati, Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. "Every Picture Tells a Story: Generating Sentences from Images." In *Computer Vision—ECCV 2010*, edited by K. Daniilidis, P. Maragos, and N. Paragios. Lecture Notes in Computer Science, Vol. 6314. Berlin: Springer.
- Fermüller, Cornelia, Fang Wang, Yezhou Yang, Konstantinos Zampogiannis, Yi Zhang, Francisco Barranco, and Michael Pfeiffer. 2018. "Prediction of Manipulation Actions." *International Journal of Computer Vision* 126 (2–4): 358–374. doi:10.1007/s11263-017-0992-z.
- Flanders, Martha, Linda Daghestani, and Alain Berthoz. 1999. "Reaching beyond Reach." *Experimental Brain Research* 126 (1): 19–30.
- Forsyth, D. A., Tamara Berg, Cecilia Ovesdotter Alm, Ali Farhadi, Julia Hockenmaier, Nicolas Loeff, and Gang Wang. 2009. "Words and Pictures: Categories, Modifiers, Depiction, and Iconography." In *Object Categorization: Computer and Human Vision Perspectives*, edited by Sven J. Dickinson, Aleš Leonardis, Bernt Schiele, and Michael J. Tarr, 167–181. Cambridge: Cambridge University Press.
- Friston, Karl, Jean Daunizeau, and Stefan Kiebel. 2009. "Reinforcement Learning or Active Inference?" *PLoS One* 4:e6421. doi:10.1371/journal.pone.0006421.

- Gandolfo, F., Giulio Sandini, and Emilio Bizzi. 1996. "A Field-Based Approach to Visuo-motor Coordination." Paper presented at the Workshop on Sensorimotor Coordination: Amphibians, Models, and Comparative Studies, Sedona, AZ.
- George, N., and Conty, L. 2008. "Facing the Gaze of Others." *Clinical Neurophysiology* 38 (3): 197–207. doi:10.1016/j.neucli.2008.03.001.
- Goodale, Melvyn A., and A. David Milner. 1992. "Separate Visual Pathways for Perception and Action." *Trends in Neurosciences* 15 (1): 20–25. doi:10.1016/0166-2236(92)90344-8.
- Gori, Monica, Alessandra Sciutti, David Burr, and Giulio Sandini. 2011. "Direct and Indirect Haptic Calibration of Visual Size Judgments." *PLoS One* 6 (10): e25599. doi:10.1371/journal.pone.0025599.
- Graziano, Michael S. A., Charlotte S. R. Taylor, Tirin Moore, and Dylan F. Cooke. 2002. "The Cortical Control of Movement Revisited." *Neuron* 36 (3): 349–362. doi:10.1016/S0896-6273(02)01003-6.
- Graziano, M. S., G. S. Yap, and C. G. Gross. 1994. "Coding of Visual Space by Premotor Neurons." *Science* 266 (5187):1054–1057. doi:10.1126/science.7973661.
- Gredebäck, Gustaf, and Terje Falck-Ytter. 2015. "Eye Movements during Action Observation." *Perspectives on Psychological Science: A Journal of the Association for Psychological Science* 10 (5): 591–598. doi:10.1177/1745691615589103.
- Gu, Chunhui, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, et al. 2018. "Ava: A Video Dataset of Spatio-temporally Localized Atomic Visual Actions." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6047–6056. New York: IEEE.
- Gupta, Abhinav, Aniruddha Kembhavi, and Larry Davis. 2009. "Observing Human-Object Interactions: Using Spatial and Functional Compatibility for Recognition." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31:1775–1789. doi:10.1109/tpami.2009.83.
- Hall, Joanna, Terry Tritton, Angela Rowe, Anthony Pipe, Chris Melhuish, and Ute Leonards. 2014. "Perception of Own and Robot Engagement in Human–Robot Interactions and Their Dependence on Robotics Knowledge." *Robotics and Autonomous Systems* 62 (3): 392–399. <https://doi.org/10.1016/j.robot.2013.09.012>.
- Hamilton, Antonia, Dan W. Joyce, J. Robert Flanagan, Chris D. Frith, and Daniel M. Wolpert. 2007. "Kinematic Cues in Perceptual Weight Judgement and Their Origins in Box Lifting." *Psychological Research* 71 (1): 13–21.
- Hamilton, Antonia, Daniel Wolpert, and Uta Frith. 2004. "Your Own Action Influences How You Perceive Another Person's Action." *Current Biology* 14 (6): 493–498. doi:10.1016/j.cub.2004.03.007.
- Held, Richard, and Alan Hein. 1963. "Movement-Produced Stimulation in the Development of Visually Guided Behavior." *Journal of Comparative and Physiological Psychology* 56 (5): 872.
- Jamone, Lorenzo, Emre Ugur, Angelo Cangelosi, Luciano Fadiga, Alexandre Bernardino, Justus Piater, and José Santos-Victor. 2016. "Affordances in Psychology, Neuroscience, and Robotics: A Survey." *IEEE Transactions on Cognitive and Developmental Systems* 10 (1): 4–25. doi:10.1109/tcds.2016.2594134.
- Jeannerod, Marc. 2001. "Neural Simulation of Action: A Unifying Mechanism for Motor Cognition." *Neuroimage* 14:S103–109. doi:10.1006/nimg.2001.0832.
- Jeannerod, Marc, and Pierre Jacob. 2005. "Visual Cognition: A New Look at the Two-Visual Systems Model." *Neuropsychologia* 43 (2): 301–312. doi:10.1016/j.neuropsychologia.2004.11.016.
- Johansson, Roland S., Göran Westling, Anders Bäckström, and J. Randall Flanagan. 2001. "Eye-Hand Coordination in Object Manipulation." *Journal of Neuroscience* 21 (17): 6917–6932.
- Kotseruba, Iuliia, and John K. Tsotsos. 2020. "40 Years of Cognitive Architectures: Core Cognitive Abilities and Practical Applications." *Artificial Intelligence Review* 53 (1): 17–94. doi:10.1007/s10462-018-9646-y.
- Lacquaniti, Francesco, Carlo Terzuolo, and Paolo Viviani. 1983. "The Law Relating the Kinematic and Figural Aspects of Drawing Movements." *Acta Psychologica* 54 (1):115–130. [https://doi.org/10.1016/0001-6918\(83\)90027-6](https://doi.org/10.1016/0001-6918(83)90027-6).
- Laird, John. 2012. *The Soar Cognitive Architecture*. Cambridge, MA: MIT Press.
- Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Miller, George, and Christiane Fellbaum. 2007. "Wordnet Then and Now." *Language Resources and Evaluation* 41:209–214. doi:10.1007/s10579-007-9044-6.
- Milner, David, and Mel Goodale. 1995. *The Visual Brain in Action*. Vol. 27. Oxford: Oxford University Press.
- Mishra, Ajay, and Yiannis Aloimonos. 2009. "Active Segmentation." *International Journal of Humanoid Robotics* 6:361–386. doi:10.1142/s0219843609001784.
- Mishra, Ajay, Yiannis Aloimonos, Loong-Fah Cheong, and Ashraf Kassim. 2012. "Active Visual Segmentation." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34:639–53. doi:10.1109/tpami.2011.171.
- Morasso, Pietro, and Ferdinando A. Mussa Ivaldi. 1982. "Trajectory Formation and Handwriting: A Computational Model." *Biological Cybernetics* 45 (2):131–142. doi:10.1007/bf00335240.

- Myers, Austin, Ching Teo, Cornelia Fermüller, and Yiannis Aloimonos. 2015. "Affordance Detection of Tool Parts from Geometric Features." *Proceedings—IEEE International Conference on Robotics and Automation*, 1374–1381. New York: IEEE. doi:10.1109/icra.2015.7139369.
- Nagai, Yukie. 2019. "Predictive Learning: Its Key Role in Early Cognitive Development." *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences* 374 (1771): 20180030. doi: 10.1098/rstb.2018.0030.
- Ogale, Abhijit S., and Yiannis Aloimonos. 2007. "A Roadmap to the Integration of Early Visual Modules." *International Journal of Computer Vision* 72 (1): 9–25. doi:10.1007/s11263-006-8890-9.
- Palinko, Oskar, Francesco Rea, Giulio Sandini, and Alessandra Sciutti. 2015. "Eye Gaze Tracking for a Humanoid Robot." In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 318–324. New York: IEEE.
- Palinko, Oskar, Francesco Rea, Giulio Sandini, and Alessandra Sciutti. 2016. "Robot Reading Human Gaze: Why Eye Tracking Is Better than Head Tracking for Human-Robot Collaboration." In *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5048–5054. New York: IEEE.
- Palinko, Oskar, Alessandra Sciutti, Lars Schillingmann, Francesco Rea, Yukie Nagai, and Giulio Sandini. 2015. "Gaze Contingency in Turn-Taking for Human Robot Interaction: Advantages and Drawbacks." In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication*, 369–374. New York: IEEE.
- Pastra, Katerina, and Yiannis Aloimonos. 2012. "The Minimalist Grammar of Action." *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences* 367:103–117. doi:10.1098/rstb.2011.0123.
- Pastra, Katerina, Eirini Balta, Panagiotis Dimitrakis, and Giorgos Karakatsiotis. 2011. "Embodied Language Processing: A New Generation of Language Technology." *Language-Action Tools for Cognitive Artificial Agents* 11:14.
- Poeticon Project. 2012. <http://www.poeticon.eu>.
- Rao, Rajesh, and Dana Ballard. 1997. "Dynamic Model of Visual Recognition Predicts Neural Response Properties in the Visual Cortex." *Neural Computation* 9:721–763. doi:10.1162/neco.1997.9.4.721.
- Regolin, Lucia, Luca Tommasi, and Giorgio Vallortigara. 2000. "Visual Perception of Biological Motion in Newly Hatched Chicks as Revealed by an Imprinting Procedure." *Animal Cognition* 3 (1): 53–60. doi:10.1007/s100710050050.
- Richardson, Magnus J. E., and Tamar Flash. 2002. "Comparing Smooth Arm Movements with the Two-Thirds Power Law and the Related Segmented-Control Hypothesis." *Journal of Neuroscience* 22 (18): 8201. doi:10.1523/jneurosci.22-18-08201.
- Sandini, Giulio. 1997. "Artificial Systems and Neuroscience." Otto and Martha Fischbeck Seminar on Active Vision, April, Berlin.
- Sandini, Giulio, F. Gandolfo, Enrico Grosso, and Massimo Tristarelli. 1993. "Vision during Action." In Vol. 8, *Active Perception*, edited by John Aloimonos, 292. Hillsdale, NJ: Lawrence Erlbaum.
- Sandini, Giulio, Giorgio Metta, and Juergen Konczak. 1997. "Human Sensori-motor Development and Artificial Systems." Paper presented at the International Symposium on Artificial Intelligence, Robotics, and Intellectual Human Activity Support for Nuclear Applications, Japan.
- Sandini, Giulio, Alessandra Sciutti, and Francesco Rea. 2019. "Movement-Based Communication for Humanoid-Human Interaction." In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 2169–2197. Dordrecht: Springer Netherlands.
- Sciutti, Alessandra, Caterina Ansuini, Cristina Becchio, and Giulio Sandini. 2015. "Investigating the Ability to Read Others' Intentions Using Humanoid Robots." *Frontiers in Psychology* 6:1362. doi:10.3389/fpsyg.2015.01362.
- Sciutti, Alessandra, A. Bisio, F. Nori, G. Metta, L. Fadiga, T. Pozzo, and G. Sandini. 2012. "Measuring Human-Robot Interaction through Motor Resonance." *International Journal of Social Robotics* 4 (3):223–234.
- Sciutti, Alessandra, Martina Mara, Vincenzo Tagliasco, and Giulio Sandini. 2018. "Humanizing Human-Robot Interaction: On the Importance of Mutual Understanding." *IEEE Technology and Society Magazine* 37 (1): 22–29.
- Sciutti, Alessandra, Laura Patanè, Francesco Nori, and Giulio Sandini. 2014. "Understanding Object Weight from Human and Humanoid Lifting Actions." *IEEE Transactions on Autonomous Mental Development* 6 (2): 80–92. doi:10.1109/tamd.2014.2312399.
- Sciutti, Alessandra, Laura Patanè, and Giulio Sandini. 2019. "Development of Visual Perception of Others' Actions: Children's Judgment of Lifted Weight." *PLoS One* 14 (11): e0224979. doi:10.1371/journal.pone.0224979.
- Senot, Patrice, Alessandro D'Ausilio, Michele Franca, Luana Caselli, Laila Craighero, and Luciano Fadiga. 2011. "Effect of Weight-Related Labels on Corticospinal Excitability during Observation of Grasping: A TMS Study." *Experimental Brain Research* 211 (1): 161–167. doi:10.1007/s00221-011-2635-x.

- Sigurdsson, Gunnar A., Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. 2016. "Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding." In *European Conference on Computer Vision*, 510–526. Cham, Switzerland: Springer.
- Soomro, Khurram, Amir Roshan Zamir, and Mubarak Shah. 2012. "Ucf101: A Dataset of 101 Human Action Classes from Videos in the Wild." ArXiv preprint: 1212.0402.
- Stein, Barry E., and M. Alex Meredith. 1993. *The Merging of the Senses*. Cognitive Neuroscience Series. Cambridge, MA: MIT Press.
- Summers-Stay, Douglas, Ching L. Teo, Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos. 2012. "Using a Minimal Action Grammar for Activity Understanding in the Real World." In *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4104–4111. New York: IEEE.
- Tani, Jun, and Stefano Nolfi. 1999. "Learning to Perceive the World as Articulated: An Approach for Hierarchical Learning in Sensory-Motor Systems." *Neural Networks* 12:1131–1141. doi:10.1016/s0893-6080(99)00060-x.
- Teo, Ching L., Yezhou Yang, Hal Daumé, Cornelia Fermüller, and Yiannis Aloimonos. 2012. "Towards a Watson That Sees: Language-Guided Action Recognition for Robots." In *2012 IEEE International Conference on Robotics and Automation*, 374–381. New York: IEEE.
- Tomasello, Michael, Brian Hare, Hagen Lehmann, and Josep Call. 2007. "Reliance on Head versus Eyes in the Gaze Following of Great Apes and Human Infants: The Cooperative Eye Hypothesis." *Journal of Human Evolution* 52 (3): 314–320. doi:10.1016/j.jhevol.2006.10.001.
- Ungerleider, Leslie G., and M. Mishkin. 1982. "Two Cortical Visual Systems." In *Analysis of Visual Behavior*, edited by M. A. Goodale, D. J. Ingle, and R. J. W. Mansfield, 549–586. Cambridge, MA: MIT Press.
- Vannucci, Fabio, G. Di Cesare, Francesco Rea, Giulio Sandini, and Alessandra Sciutti. 2018. "A Robot with Style: Can Robotic Attitudes Influence Human Actions?" In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, 1–6. New York: IEEE.
- Vernon, David. 2006. "The Space of Cognitive Vision." In *Cognitive Vision Systems*, edited by H. I. Christensen and H. H. Nagel, 7–24. Berlin: Springer-Verlag.
- Vernon, David, Claes von Hofsten, and Luciano Fadiga. 2010. *A Roadmap for Cognitive Development in Humanoid Robots*. Cognitive Systems Monographs. Berlin: Springer.
- Verschure, Paul. 2012. "Distributed Adaptive Control: A Theory of the Mind, Brain, Body Nexus." *Biologically Inspired Cognitive Architectures* 1:55–72. doi:10.1016/j.bica.2012.04.005.
- Vignolo, Alessia, Nicoletta Noceti, Francesco Rea, Alessandra Sciutti, Francesca Odone, and Giulio Sandini. 2017. "Detecting Biological Motion for Human–Robot Interaction: A Link between Perception and Action." *Frontiers in Robotics and Ai* 4:14.
- Viviani, Paolo, and Carlo Terzuolo. 1982. "Trajectory Determines Movement Dynamics." *Neuroscience* 7 (2): 431–437. doi:10.1016/0306-4522(82)90277-9.
- Vondrick, Carl, Hamed Pirsiavash, and Antonio Torralba. 2016. "Anticipating Visual Representations from Unlabeled Video." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 98–106. New York: IEEE.
- Wang, Zhidan, Rebecca A. Williamson, and Andrew N. Meltzoff. 2018. "Preschool Physics: Using the Invisible Property of Weight in Causal Reasoning Tasks." *PLoS One* 13 (3): e0192054. doi:10.1371/journal.pone.0192054.
- Wu, Ying, and Thomas S. Huang. 1999. "Vision-Based Gesture Recognition: A Review." In *International Gesture Workshop*, 103–115. Berlin: Springer.
- Xu, Kelvin, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention." *Proceedings of the 32nd International Conference on Machine Learning*, PMLR 37:2048–2057.
- Yang, Yezhou, Anupam Guha, Cornelia Fermüller, and Yiannis Aloimonos. 2014. "A Cognitive System for Understanding Human Manipulation Actions." *Advances in Cognitive Systems* 3:67–86.
- Yu, Xiaodong, Cornelia Fermüller, Ching Lik Teo, Yezhou Yang, and Yiannis Aloimonos. 2011. "Active Scene Recognition with Vision and Language." In *2011 International Conference on Computer Vision*, 810–817. New York: IEEE.

15 Cognitive Robot Navigation

Jiru Wang, Jianxin Peng, Rui Yan, and Huajin Tang

15.1 Introduction

Navigation in cognitive robotics has been strongly influenced by studies on navigation in animals. During the later decades of the twentieth century, researchers have focused on studying rats' spatial learning and memory in mazes to help understand the idea of spatial cognition for other species, including humans. The fact that a rat can reach flexible target locations effortlessly in complex mazes inspired scientists to determine that the flexible movement behavior is dependent on an inner map formed in the brain. This inner map can reflect the spatial and geometric relations between animals and surroundings. Furthermore, by observing animals' behaviors, including rats, bats, and more, researchers obtained two important findings: 1) animals were able to successfully return home even when put into a seldomly visited place; 2) animals looked for shortcuts. The two findings could verify that the inner map made it possible for the evaluation of relative positions and navigation from the current position to target places.

By behavioral observation and psychological analysis, researchers started to study animals' spatial cognition (Tolman 1948) and to understand the spatial cognition abilities in complex environments. The concept of a "cognitive map," proposed by Tolman in 1948, has been widely considered to possibly provide the basis for spatial memory and navigation. In order to reveal how animals construct cognitive maps of environments, studies in arthropods found the existence of a highly effective path integration mechanism depending on directional heading and distance computations. Then important discoveries about spatial cells in mammals demonstrated that the path integration mechanism completed by some brain regions is necessary to form inner cognitive maps. These maps represent the topological structures of environments and surrounding landmarks by position coordinates. With the discoveries of place cells (O'Keefe and Dostrovsky 1971), head direction cells (Taube, Muller, and Ranck 1990), and grid cells (Hafting et al. 2005), neuroscientists began to study the mechanisms underlying spatial navigation skills in animals. This research became a milestone of cognitive map and spatial navigation research.

The cognitive map mechanism of animals provides good insight to develop bioinspired models of spatial cognition for robots. Animals can perform simultaneous localization and mapping (SLAM) robustly and effortlessly in daily life. They can also quickly adapt to

new dynamic environments and localize themselves. Based on psychological and neuroscientific studies on animal spatial navigation, researchers have attempted to create bio-inspired map building simulations and make the spatial navigation of mobile robots more flexible and robust (Milford, Wyeth, and Prasser 2004). The target is to make more stable and general intelligent navigation systems for robots to increase the capabilities of autonomy and operation flexibility.

15.2 From Psychology to Neuroscience

In the 1930s, E. C. Tolman started to research cognitive behavioral psychology by observing rats running in various types of mazes. Experiments showed that rats could plan paths with fewer and fewer mistakes until they finally completed the correct path planning. Thus, Tolman concluded that there is one kind of inner mental knowledge structure in an animal's brain that stores information according to the animal's position in the environment. Tolman (1948) then proposed the concept of the cognitive map in 1948. The key findings of the cognitive map include latent learning and spatial learning. Latent learning means that the rats learn about the structure of the maze without getting a food reward and can quickly plan the optimal path in the maze once food is given. And in the sunburst maze, the rats first learn to plan specific paths according to different rewards. If the planned path is blocked, they can still find an optimal path they have not previously experienced. This ability has been called spatial learning. The cognitive map theory directly sets the stage for studies about how space is represented in the brain.

Neurophysiological experiments have helped to verify Tolman's cognitive map theory and have searched for the neural basis of the cognitive map mechanism in animals' brains from neural structures and cell activities. In one such prominent and successful experiment, Hubel et al. (Hubel and Wiesel 1959, 1977) inserted electrodes into specific brain regions of awake animals and were able to observe and record neurons' activities in the cortex. They demonstrated that animals' specific functional behavior can be understood by neural activity and the interaction between neurons. After that, extensive neuroscience research studies at the neuron and synapse level were carried out to establish the relationship of synaptic physiology and animal behavior. Based on the advancement of neurophysiological experimental techniques, early research on the hippocampal region provided good insight and resulted in widespread agreement that the hippocampus is an important region for encoding and maintaining memories. In another set of studies, neuroscientists were motivated to associate specific individual behaviors with neuronal population activities in specific brain regions, including the hippocampus. A series of studies about special firing patterns of cells in the hippocampus and the surrounding regions were performed to unveil the spatial cognition and navigation mechanism in animal brains. In these studies, when rats moved into controlled maze environments, the activity of a single neuron or a neural population (mainly in the hippocampus and surrounding areas) was recorded through an electrode inserted into a specific brain region. The accumulated experimental results led to the discoveries of multiple types of cells responsible for inner cognitive mapping. Every type of cell shows specific firing patterns for encoding the environment and thus plays an important role in animals' spatial cognition.

Place cells were discovered by O’Keefe and colleagues in the 1970s (O’Keefe and Dostrovsky 1971). These cells are in the hippocampus and fire consistently when a rat is at a particular location in the environment. The firing cell signals recognition of a specific place in an environment, known as the cell’s “place field.” It is suggested that thousands of place cells, covering the surface of any space, act as a mapping system in the hippocampus to create a cognitive map (O’Keefe and Nadel 1978). Each place cell receives two different inputs, one external input about environmental stimuli and external events and an internal input from an inner-path integration system based on its self-motion.

In the 1980s, Ranck (1984) observed strong directional tuning when cellular activity was recorded from the pre- and parasubiculum regions. These direction-tuned cells were also discovered in other brain regions (Taube 2007), such as the medial entorhinal cortex (MEC; Sargolini et al. 2006). These cells respond to an animal’s head direction and are called head direction cells. They only fire when the rat’s head is at specific orientations. All orientations are represented by the head direction cell population. About thirty years after the discovery of the place cell, grid cells were discovered in the entorhinal cortex (EC) by Edvard I. Moser (Fyhn et al. 2004). Grid cells show the properties most like place cells but have multiple firing fields (figure 15.1)—that is, they can fire in a metrically regular way on the whole surface of a given environment. The firing fields of these cells have been demonstrated to be in a hexagonal pattern. In fact, a single grid cell will fire when the rat is located at any of the vertices of a tessellating hexagonal pattern. Grid cell firing appears to be a signal used for measuring displacement distances and direction—in other words, a “metric.” Grid cells differ from each other in grid spacing, phase, and orientation (Hafting et al. 2005; Fyhn et al. 2004), and the spacing of grid cells increases along the dorsal-ventral axis of the EC (Brun et al. 2008).

In the same parahippocampal brain regions are additional cells, called border cells, related to spatial mapping. The border cells can achieve responses when the animal is near a boundary of the local environment (Solstad et al. 2008; Savelli, Yoganarasimha, and Knierim 2008). Boundary-related cells have also been recorded in the subiculum, which




Figure 15.1

(a) The path on which the rat traveled in a square maze and the firing of a grid cell. (b) The firing rate of the grid cell at each place. *Source:* Moser and Moser 2007.

indirectly links the feedback from CA1 to the MEC, the presubiculum and parasubiculum (Lever et al. 2009).

Neuroscientific experiments show us a number of neural representations of the inner cognitive map. They might have innate connection circuitries and together could constitute a metric navigation system: head direction cells are responsible for direction tuning; grid cells play the important roles in path integration; border cells are used for evaluating vicinity to boundaries, and place cells are taken as the place representation. The discovery of spatial cells made the cognitive map theory more dominant in spatial cognition research.

15.3 Computational Theories on Robot Spatial Cognition

15.3.1 Path Integration

Path integration means to estimate positions and plan paths to targets via the continuous integration of movement cues such as directional heading and distance over the whole path. Inspired by animal behaviors, head direction cells are responsible for orientation tuning, grid cells can execute path integration, and place cells contribute to representing places. In order to build cognitive maps, outputs from head direction cells are first considered as the input signals for grid cells, then place cells and grid cells provide a population-encoding method for path integration. Most researchers have reached a consensus on this topic, but a few important questions still rise: How do we simulate the direction-tuning characteristic of head direction cells? How do we provide grid-cell-encoding methods for path integration? How do we represent place cells using grid cells?

Direction tuning

Information processing in biological systems is generally considered to be nonlinear dynamic and can be implemented by neural networks. Stable, persistent activity has been thought important for neural computation. Amit (1989) suggested that persistent neural activity in biological networks is a result of dynamical attractors in the state space of recurrent biological networks. This study resulted in the increasing popularity of using attractor networks in neuroscientific simulation and biologically inspired system building. In addition, there was evidence that many brain areas act as attractor networks (Wills et al. 2005), including the hippocampus and the entorhinal cortex. Because of the association with the ability to continuously track changing stimuli in certain brain regions, continuous attractor dynamics are widely used for brain mechanism simulation (McNaughton et al. 1996; Trappenberg 2002).

Simulations of head direction cells can be organized in a ring attractor and modeled as a one-dimensional continuous attractor network (CAN). In this network, the head's angular velocity (inner signals provided by other brain regions) is integrated for head direction representation updating. The rat's turning range (360°) is evenly discretized into N parts that correspond to N neurons, and each neuron has weighted connections to others, as shown in figure 15.2 (Skaggs et al. 1994). The connection strength decreases with increasing distance between neurons and active neurons, and then only one direction is focused at each time point. Activity in one part of the ring is initialized by visual input from visual




Figure 15.2
One-dimensional CAN modeling head direction cell. *Source:* Skaggs et al. 1994.

cells. When the animal turns its head, sensory inputs (mainly from the vestibular region) can detect the change that activates rotation cells and cause activity bumps to move in the appropriate direction around the ring, keeping the representation concordant with the real head direction (Calton et al. 2008; Knierim and Zhang 2012).

Path integration

Currently, some proposed computational models of grid cells include oscillatory interference (OI) models (Burgess 2008; Zilli and Hasselmo 2010) and CAN models (Fuhs and Touretzky 2006; Burak and Fiete 2009). In OI models, the grid pattern arises from several oscillators with slightly different frequencies around the theta frequency (Blair, Welday, and Zhang 2007; Burgess, Barry, and O’Keefe 2007). The key requirement is that the frequency is modulated by the animal’s velocity. Under appropriate conditions, the beat frequencies of the interference patterns cause a cell to reach its firing threshold whenever the rat is at the grid vertex. The key assumptions of oscillatory models have been experimentally challenged because theta oscillations have not been observed in fruit bats (Yartsev, Witter, and Ulanovsky 2011) and macaque monkeys (Killian, Jutras, and Buffalo 2012), despite robust grid cell activity having been recorded in both species.

A grid cell model with a single grid scale based on CAN models was proposed (Burak and Fiete 2009) to perform path integration with noise-free velocity inputs. Based on different spacing in grid cells, a grid cell model with multiple grid scales is required for the path integration. In this case, neurons are often arranged in a two-dimensional neural sheet. Recurrent connectivity among neurons with global inhibition leads to random patterns of population activity that spontaneously merge into organized “bumps” of grid cell population activity. A response from the grid cell can be obtained by accumulating the firing activity of a single neuron over a full trajectory. The most remarkable progress in the field (Burak and Fiete 2009) has been to accurately integrate velocity inputs into grid cell models.

From grid cell to place cell

Functionally, path integration can be accomplished by grid cells driven by the rat's moving velocities in particular directions. Anatomically, the majority of the principal cells in layers II and III of the MEC have grid properties (Sargolini et al. 2006), and there is a strong projection from the MEC to the hippocampus. Therefore, place cell activities might have been, at least partially, generated in response to stimulation from grid cells. As the size and spacing of grid patterns increase from small in the dorsal MEC to large in the ventral MEC (Fyhn et al. 2004; Hafting et al. 2005), it is believed that the input for place cells comes from a combination of several grid cells. Linear combinations of grid fields are generally used for generating firing fields of place cells (O'Keefe and Burgess 2005; Hafting et al. 2005; McNaughton et al. 2006; Solstad, Moser, and Einvoll 2006). Although other mechanisms, such as feedback inhibition of place cells, can achieve similar results (Monaco and Abbott 2011), linear-combination-based models are easy to implement and widely employed. As each place cell receives a subset of grid cells as its input afferent, a learning algorithm is required to do selection. Since Hebbian learning is commonly accepted as a biologically plausible theory for synaptic adaptation, it was chosen to determine the connection weights between the place cell and grid cell input (Hu et al. 2016). Furthermore, the learning performances of different variations of Hebbian learning have also been compared, and potential mechanisms to improve the learning process have been discussed.

Different learning rules have been tested, and the experimental results are shown in figure 15.3 (Hu et al. 2016). The presynaptically gated learning shows better results with fewer bumps. The gated input stimulation removes the enhancement of unnecessary inputs from grid cells. The introduction of a spatial-learning window weakens stimulation from unnecessary afferents and enhances certain inputs that contribute most to place cells. Hebbian learning refines the place cell activity to fewer bumps, but place cells tend to have more and more bumps during learning without a mechanism to prevent multiple bumps. Therefore, a circle-shaped learning window is applied to the learning process so the number of bumps can be reduced. As shown in figure 15.3, only two bumps are left with the help of the spatial-learning window. To further explore the effect of learning, synaptic weights are examined from grid cells to place cells after learning. Initially, synaptic weights from grid cells to place cells are randomized with a normal distribution. As learning proceeds, synaptic weights from contributing grid cells to corresponding receiving place cells are enhanced. After learning, each place cell is expected to be strongly connected to a subset of grid cells.

One should notice that the current network structure has been simplified. In the current setting, grid cells in the same neuron sheet share the same synaptic weights as place cells. One should notice that stimulation from grid cells provides only partial information for place cells. Other sensory information, such as visual, auditory, and olfactory signals, may also affect the learning process.

15.3.2 Cognitive Map Building

Evidence has revealed that rats can correct accumulative movement errors in path integration when they meet salient landmarks (McNaughton et al. 2006). When a rat returns to a familiar environment, the path integrator should be reset to adjust to the perceived environment (Moser, Kropff, and Moser 2008). However, it remains unclear how the brain

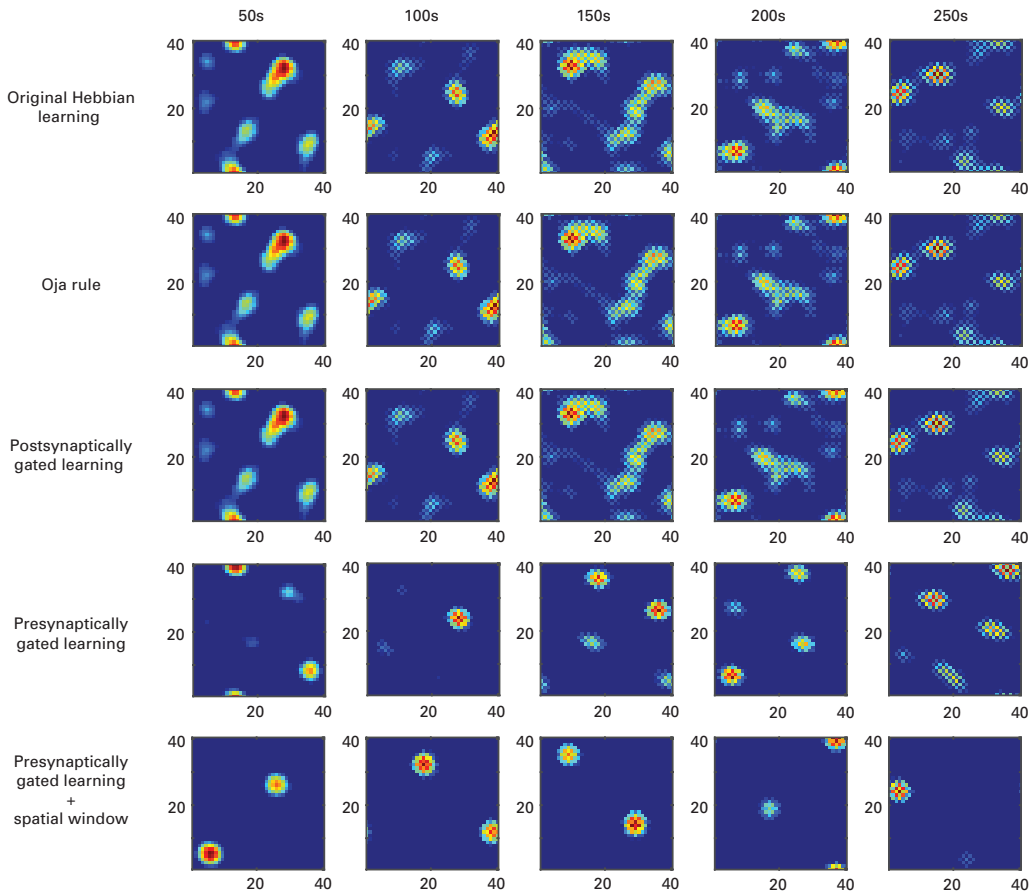


Figure 15.3
Neural activities of place cells with different learning algorithms. *Source:* Hu et al. 2016.

senses and transforms external sensory inputs into an internal cognitive map (Burak and Fiete 2009). The cognitive map theory has continuously inspired important advances in robotic mapping and navigation. The multimodal integration of visual place cells and grid cells has been proposed to enhance robot localization (Cuperlier, Quoy, and Gaussier 2007; Jauffret et al. 2012).

Milford et al. made significant progress in emulating the spatial navigation ability of the hippocampal system by building a semimetric topological map in a very large area (Milford and Wyeth 2008, 2010). In their work, the core model, RatSLAM, has been demonstrated to build maps simultaneously in large and complex environments. It emulates the rat's spatial-encoding behavior using three key components: the pose cells that are analogous to the rodent's conjunctive grid cells, the local view cells that provide the interface to the robot's sensors in place of the rodent's perceptual system, and the experience map that functionally replaces place cells. Each local view cell is associated with a distinct visual scene of the environment and activated when the robot sees that scene. A CAN is built for pose cells to encode the estimate of the robot's pose. Each pose cell is connected to proximal cells by excitatory and inhibitory connections with wrapping across

all six faces of the network. Intermediate layers in the xy plane are not shown. The network connectivity leads to clusters of active cells known as activity packets. Active local view and pose cells drive the creation of experience nodes in the experience map, a semimetric graphic representation of visited places in the environment and their interconnectivity. RatSLAM is an attempt to build a practical robotic system to take advantage of the spatial navigation mechanism highlighted by studies of the rat brain. It can perform well for some challenging problems in robotic navigation. The maps based on RatSLAM are less accurate than those of traditional SLAM systems, but its flexibility can help to cope with noisy input, deal with a changing environment, and accommodate increasing complexity. The cognitive map building mechanism in bioinspired SLAM will create a new generation of lightweight and low-cost mapping and navigation systems to be deployed in the robotic navigation field for real and large environments.

In 2013, Steckel and Peremans (2013) proposed a biomimetic navigation model named BatSLAM, which can solve simultaneous localization and map tasks with a biomimetic sonar mounted on a mobile robot. They analyzed the performance of the proposed robotic implementation operating in the real world and concluded that the biomimetic navigation model operating on the information from the biomimetic sonar can allow an autonomous agent to map unmodified environments efficiently and consistently. This showed that consistent topological maps with semimetric properties can be constructed using only motor commands and biomimetic sonar “fingerprints.” Furthermore, if these sonar “fingerprints” are sufficiently informative, there is no requirement for further interpretation of them in terms of discrete objects positioned in the environment.

In 2015, Silveira et al. (2015) presented a new bioinspired algorithm for underwater SLAM called DolphinSLAM, which extended the successful previous RatSLAM approach from 2D ground vehicles to 3D underwater environments. The proposed model uses a neural network model to localize and deal with low-resolution monocular images and imaging sonar data, in contrast to other available navigation systems that focus on probabilistic methods and occupancy grids. The model is composed of six modules: the preprocessing module, the local view recognition module, the motion detection module, the 3D place cells network module, and an experience map module. It has the particular advantage of being an appearance-based navigation system that can work well with low-resolution sonar and visual image data, in contrast to other available navigation underwater systems that focus on probabilistic methods.

Together with sensory-information processing, grid and place cells are considered to afford animals with an innate sense of the world around them. Inspired by the path integration mechanism of grid cells, Yuan et al. (2015) proposed a cognitive map model (figure 15.4), simulating grid and place cells for path integration and place representation. Visual cues are used for the error correction and cell population activity, resetting when loop closures are detected. Depth information in visual cues is invariant to lighting conditions and makes some similar indoor scenes more distinguishable. A comparison between image profiles is performed for each pair of incoming RGB and depth frames for loop closure and new scene detection. More details can be found in Tian et al. (2013).

In this work the cognitive map contains a set of spatial coordinates that the robot has experienced in its past travels. The robot’s spatial coordinates are calculated from place cell population activities, which are generated from a subset of grid cell population activities. Nodes in the cognitive map are constructed by associating the major peak of the place

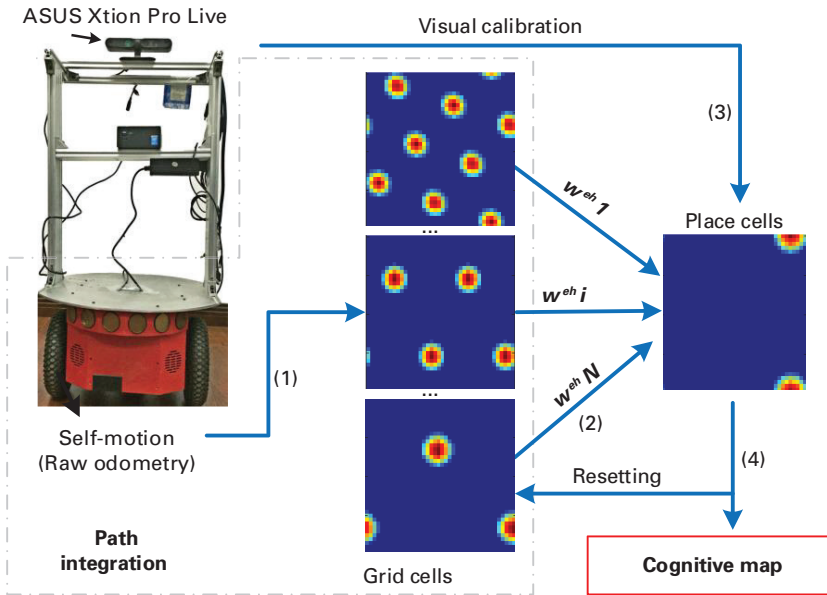


Figure 15.4

The system architecture of a cognitive map building model. *Source:* Yuan et al. 2015.

cell population activities with corresponding visual cues and locations being denoted as visual experiences. Algorithm 1, below, shows the cognitive map building process. The incoming visual inputs are compared with past visual experiences. If the latest input matches the previous visual experience, it is considered a familiar scene that the robot has previously seen. The status of the grid cell population activities and the place cell population activities is then reset to the previous matched visual experiences. The current visual input and the matched visual experience are merged into the same experience. Otherwise, a new visual experience is created. Once a loop closure is detected, the map will be adjusted to the recalled visual experiences.

Algorithm 1. The Cognitive Map Building Algorithm

Input: Raw odometry data from wheel encoders and visual images from the RGB-D sensor (1)

Output: Cognitive map

Begin: Calculate grid cell population activities (2)

 Calculate place cell population activities (3)

 Obtain one major peak of place cell population activities

 Perform visual profile comparison (4)

if the incoming visual input matches the previous visual experiences

then perform resetting and map correction (6)

else create a new visual experience (5)

end if

End

A cognitive map for a large office environment of $35\text{ m} \times 35\text{ m}$ on a mobile robot was built to validate the effectiveness of the proposed model (for more details of the parameter setting and platforms, see Yuan et al. [2015]). Figure 15.5 demonstrates the experimental results. Row A shows the dead-reckoning map obtained from the robot odometry. Obviously, this map cannot represent the environment properly. Row B shows the cognitive map based on the proposed computational model. With visual inputs, the system can successfully perform loop closure detection and correct the odometry drift. Finally, it generates a cognitive map that encodes both topological and metric information. In Row C, the

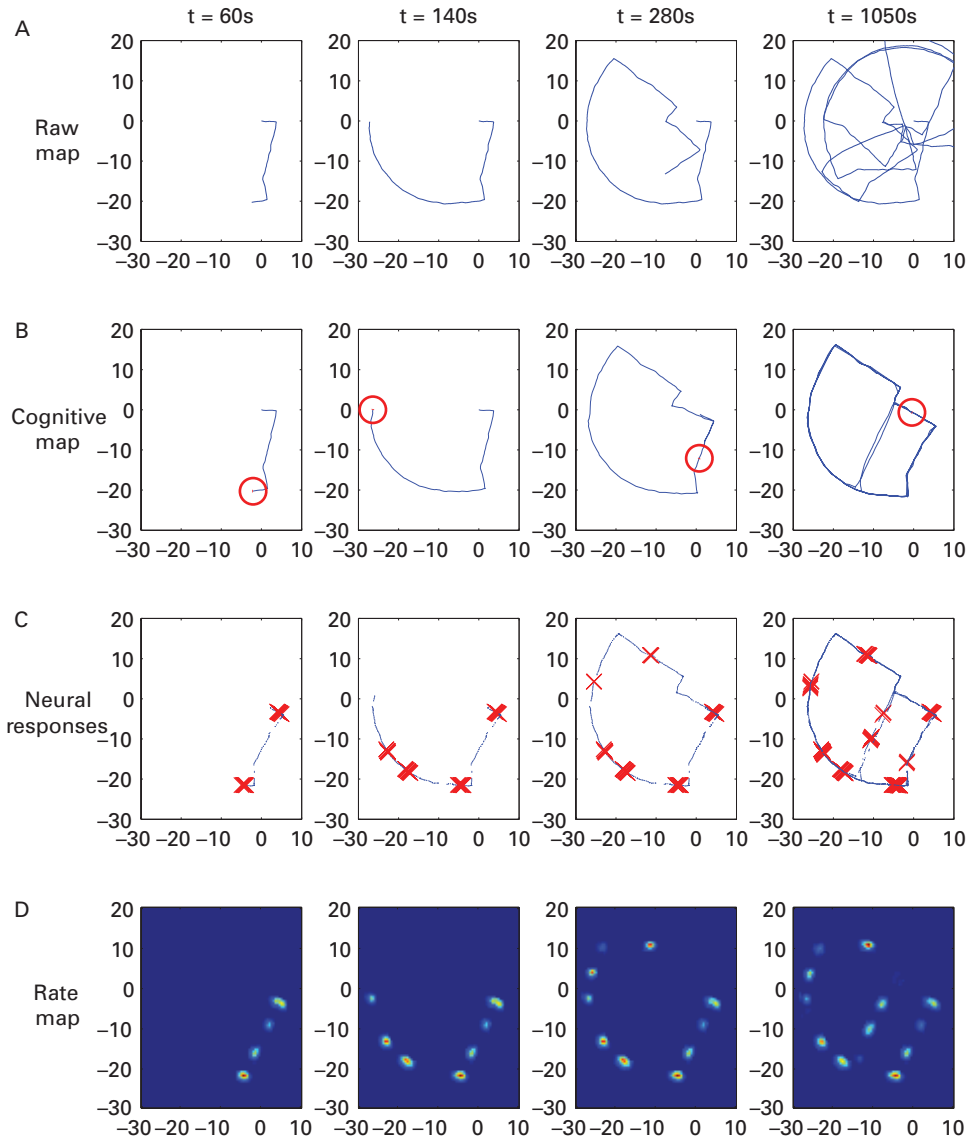


Figure 15.5
Neural responses in the map building process. *Source:* Yuan et al. 2015.

blue dotted line shows the real trajectory traveled by the robot, and the red crosses indicate the firing locations of the grid cell located at (20, 20) in the twenty-first layer of the neural sheets. Row D shows the performance of maps corresponding to different rates. To generate the rate map, a spatial smoothing algorithm with a Gaussian kernel, as described in Hafting et al. (2005), is adopted with a bin size of $0.5 \text{ m} \times 0.5 \text{ m}$.

It is a significant challenge to build robust SLAM systems in dynamical large-scale environments. Inspired by findings in the entorhinal-hippocampal neuronal circuits, Zeng and Si (2017) proposed a cognitive mapping model that includes continuous attractor networks of head-direction cells and conjunctive grid cells to integrate velocity information by conjunctive encodings of space and movement. Visual inputs from the local view cells in the model provide feedback cues to correct drifting errors of the attractors caused by the noisy velocity inputs. The key components of the proposed model include head direction (HD) cells, conjunctive grid cells, and local view cells. Both HD cells and conjunctive grid cells are modeled by continuous attractor networks that operate on the same principles. The conjunctive representations of space and movement allow the networks to reach stable states for all movement conditions. And the proposed model is robust in building a coherent semimetric topological map of the entire urban area using a monocular camera, even though the image inputs contain various changes caused by different light conditions and terrains.

Animals such as birds and bats possess superlative navigation capabilities, robustly navigating over vast three-dimensional environments and leveraging an internal neural representation of space combined with external sensory cues and self-motion cues. Yu et al. (2019) presented a novel, neuroinspired 4-DOF (degrees of freedom) SLAM system named NeuroSLAM for mapping and localization in large, real-world three-dimensional environments that integrated with a vision system that provides external visual cues and self-motion cues. In this model, the robot's state of a 4-DOF pose (x, y, z, yaw) in 3D environments is represented by the activity in the 3D grid cell network and the multilayered head direction cell network, conjunctively. The conjunctive pose cell network performs path integration on the basis of the self-motion cues and performs calibration based on the local visual cues. The approaches to the creation and relaxation of the multilayered graphic experience map are based on the combination of local view cells, conjunctive pose cells, and 3D visual odometry. The 3D multilayered experience map generated by the NeuroSLAM system can be learned and generated when the robot visits unknown environments. It can also be incrementally maintained and updated based on the learning and recalling mechanism. The 3D spatial experience nodes represent a 4-DOF pose in a specific 3D location, and the links contain distances and directions between nodes. This metric and topology information can be used for 3D path planning and guidance control in 3D environments. It is likely that map maintenance routines could also be deployed to ensure long-term map stability as well as computation and storage viability.

The computational mechanisms of mammalian brains in integrating different sensory modalities under uncertainty for navigation are enlightening for robot navigation. Zeng et al. (2020) proposed a concise yet biologically plausible model integrating visual and vestibular inputs, NeuroBayesSLAM, based on spatial cognitive mechanisms of mammalian brains to solve the SLAM problem. The proposed model successfully built coherent cognitive maps in both large-scale outdoor and small indoor environments. In the model,

the pose of the robot is encoded separately by two subnetworks—namely, a head direction network for angle representation and a grid cell network for position representation, using the similar neural codes of head direction cells and grid cells observed in mammalian brains. The neural codes in each of the subnetworks are updated in a Bayesian manner by a population of integrator cells for vestibular cue integration, as well as a population of calibration cells for visual cue calibration. The conflict between the vestibular cue and visual cue is resolved by the competitive dynamics between the two populations. The model successfully builds semimetric topological maps and self-localizes in outdoor and indoor environments with different characteristics, achieving a performance comparable to previous neurobiologically inspired navigation systems but with much less computation complexity. The proposed multisensory integration method constitutes a concise yet robust and biologically plausible method for robot navigation in large environments. The model provides a viable Bayesian mechanism for multisensory integration that may pertain to other neural subsystems beyond spatial cognition.

One should note that in most experiments (Burak and Fiete 2009; Zilli and Hasselmo 2010), velocity inputs are extracted from ground-truth trajectories. However, for animals or autonomous mobile robots, accumulated errors are inevitable. In the above model, velocity inputs were extracted from idiothetic wheel encoders to drive CAN-based grid cell population activities, and accumulated errors exist in raw odometry data. Together with visual cues for loop-closure detection and map correction, the model can produce an accurate representation of the environment and contributes to developing, innovative robotic spatial cognition approaches (Huang, Tang, and Tian 2014; Milford and Wyeth 2010), showing the potential for machines that mimic more complex activity in the brain.

15.3.3 Cognitive Navigation

Humans and animals have an instinctual ability to navigate freely in environments. However, it is a challenging task to endow a robot with this ability, as a robot needs to be integrated with several functional mechanisms, such as scene understanding, mapping, self-localization, obstacle avoidance, dead reckoning and path planning (Brooks 1999; Thrun 2002). Discoveries of spatial cells and the development of the cognitive map theory motivate researchers to use biologically plausible principles for acquiring, storing, and maintaining spatial knowledge and to explore biologically inspired navigation strategies for robots.

The use of “directions” as guidance has been raised in several studies. Méndez (2012) presented a spatial conceptual map framework to transfer cognitive human navigation behaviors to an artificial agent, which can generate route directions similar to those created by humans. This conceptual map was modeled as three levels of interconnected graphs to simulate human spatial reasoning. However, this navigation system was only tested in a simulation environment. A method for modeling environments from a route perspective was discussed in Saiki et al. (2011). The route perspective is defined as a mental tour of an environment, which is represented by a person when they are walking around the area. When describing an environment in this perspective, the terms regarding relative directions such as left and right are used. Another perspective is known as a survey perspective, which describes an environment from a top view where routes and landmarks are known in advance.

A navigation strategy considering both the route and survey perspectives, called direction-driven navigation, was presented by Shim et al. (2014). The directions extracted from a cognitive map denote the use of the survey perspective, while the execution of the directions by a mobile robot in a real environment implies the involvement of the route perspective. When traveling to a target destination, the robot is guided by a direction-driven behavior, such as following the directional guidance from someone else or from GPS, instead of closely following a global or local path.

The system architecture of the proposed navigation system is presented in figure 15.6, consisting of three main components: cognitive map building, a grid-based direction planner, and multilayered asymmetrical local navigation. The proposed grid-based direction planner (as global planner) and multilayered asymmetrical local navigation (as local planner) construct the direction-driven navigation system. The global planner plans a global path connecting its current location and the final goal destination. The local planner creates a local path, connecting the current location to a local goal destination, which follows the global path closely. Initially, images are captured by a vision sensor, and odometry is obtained from the mobile base. They are assisted by a CAN (McNaughton et al. 2006), which constructs the cognitive map of the environment.

By analyzing the constructed map, the movement directions can easily be extracted in the form of “moving forward,” “turning left,” “turning right,” and “making a U-turn.” The grid-based direction planner provides directional guidance at junctions for guiding the robot to a target location. The robot compares its current visual cues to the templates associated in the cognitive map in order to localize itself. It should be noted that the localization is crucial, as a bad localization may lead the direction planner to give wrong directions. Given a direction, the navigation system executes the corresponding action only when it conforms with real conditions. For example, the robot will not execute the “turning right” instruction when the right junction is not detected.

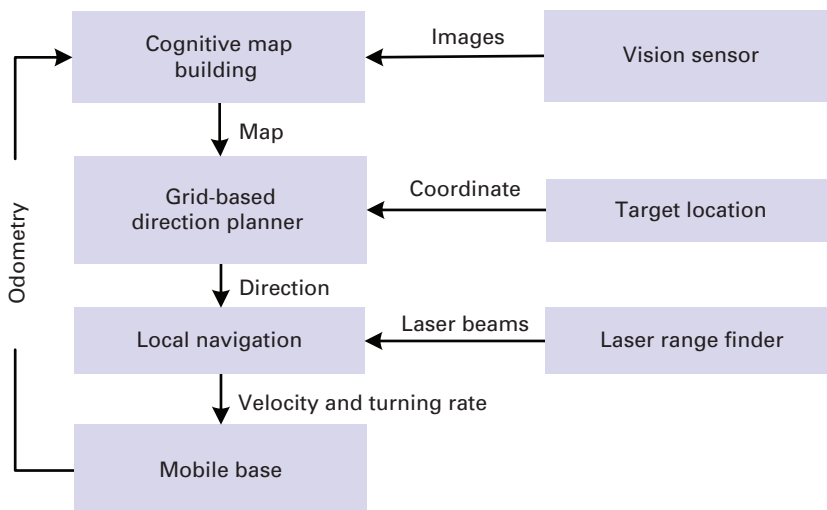


Figure 15.6

System architecture of the direction-driven navigation system. *Source:* Shim et al. 2014.

Otherwise, the proposed multilayered asymmetrical local navigation module is used to control the velocity and turning rate of the robot to guarantee a safe motion such as obstacle avoidance. A laser range finder is used as the sensor input to the local navigation module.

15.3.4 Beyond Spatial Navigation

Humans have an innate ability to explore, map, and navigate in unknown environments while simultaneously performing variant tasks. However, current technology is still far from producing a robotic servant to perform daily tasks in unstructured environments. Taking the task of serving milk tea as an example, when one orders a cup of milk tea, a robotic servant needs to understand the environment first before performing a sequence of preparatory actions at specific locations. This is a common task requiring the cognitive map and episodic memory (Buzsáki and Moser 2013), and both components play important roles for humans to perform spatiotemporal tasks. The cognitive map can provide internal spatial representations of the environment, and episodic memory for humans to learn cognitive tasks through self-experiences and then plan the actions accordingly. Biologically, the entorhinal-hippocampal region is necessary for cognitive maps and episodic memory, though it may not be sufficient (Fyhn et al. 2004; Hafting et al. 2005; Tulving and Markowitsch 1998). Functionally, the cognitive map and episodic memory form the main technologies for robotic spatial cognition. Some work has been accomplished in this field (Fleischer et al. 2007; Krichmar et al. 2005). The integration of the cognitive map and episodic memory can make the performance of the robotic system more brain-like. The cognitive map-based SLAM approaches have been successfully applied to mobile robots in real-life environments (Tian et al. 2013; Shim et al. 2014; Yuan et al. 2015). The cognitive map interfering with cognitive memory has been explored by computational modeling and applied to robotic applications (Tang et al. 2017; Hu et al. 2016).

Integrating cognitive navigation with episodic memory

Episodic memory endows humans with the ability to respond to salient events in a temporal sequence (Moser, Kropff, and Moser 2008) and recall them sequentially (Tulving and Markowitsch 1998). Though episodic memory has been studied for decades in psychology and neuroscience, recently, researchers have started to build models of episodic memory for intelligent systems. A few studies have developed episodic memory models for cognitive robots using designed data structure to simulate the functionality of episodic memory (Endo 2008; Stachowicz and Kruijff 2012; Jockel, Westhoff, and Zhang 2007). A cognitive memory network plays the role of episodic memory and is involved in navigation through recalling travel experiences, as shown in figure 15.7. This enables a robot to recognize and memorize different locations while storing and recalling the correct sequence to accomplish a task. This system takes advantage of the autoassociation of the memories through neural activities, which can achieve better flexibility and generalization abilities compared to data structure-based models relying on explicit symbolical knowledge programming. The details of the cognitive map can be found in Tang, Yan, and Tan (2018).

Episodic memory

As shown in figure 15.7, a dual network model for the CA3 region in the hippocampus is used for encoding and representing episodic memory (Tang, Yan, and Tan 2018). Both networks have synchronized gamma cycles as they share common inhibitory neurons. The

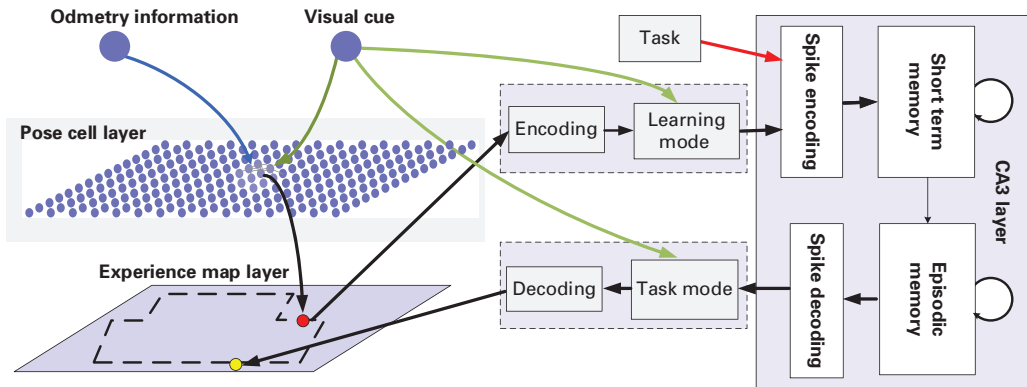


Figure 15.7

Overview of system architecture. The system is mainly divided into two parts: the cognitive map and the episodic memory. The pose information is updated by the odometry and visual input. It forms an energy package in the CAN structure. The energy package projects to the experience map, which is then converted to a grid map. In training mode, the task-related location information will be stored in memory. In task mode, the task-related location information is retrieved from memory and used to navigate the robot. *Source:* Tang, Yan, and Tan 2018.

episodic memory network stores the active sequence in its synaptic weights before transferring them to the neocortex. During recall, a cue consisting of the first two items in the desired sequence is presented to the neocortex, which will then reproduce the rest of the stored sequence. The output sequence will only be produced once and will not be repeated. The main steps will be discussed as follows:

1) **Storage:** the storage of a memory sequence is first demonstrated by introducing seven distinct items in the memory sequence to the CA3 short-term memory network. Each item is introduced to the network at the trough of the theta rhythm, and the network will repeat this value near the peak of each subsequent oscillation. Once the entire sequence has been fully introduced to the short-term memory network, the sequence is presented to the episodic memory network for storage. Here, the sequence is repeated a few times in its entirety until the episodic memory network can learn and store the sequence by updating synaptic weights. Once the storing phase is completed, the amplitude of theta rhythm is reset to zero to stop the function of short-term memory.

2) **Retrieval:** in the retrieval phase, the first two items in the memory sequence are presented directly to the neocortex as a retrieval cue. After receiving the cue, the pyramidal cells representing the first two items will fire and transmit the action potentials down through synaptic connections to subsequent memory items. Synaptic inputs from the firing of the first two memory items are sufficient to trigger the firing of the next memory item but insufficient for other items. Next, the cumulative synaptic inputs from the firing of the first three memory items trigger the firing of the fourth memory item. The process continues until the entire sequence has been triggered. Hence, the stored sequence memory is retrieved.

Exploration and navigation

The proposed architecture is verified based on a mobile robot platform Neco in a laboratory environment and a convention hall environment. The robot is equipped with sonar sensors and laser scanners for obstacle avoidance, maintaining a straight path, and detecting

turns and junctions. Neco is programmed to conduct five types of motion: moving forward, turning 90° right, turning 90° left, turning 180°, and then stopping at intersections and stopping at the end. In the task mode, after decoding the neural signals to grid indexes, the memory in CA3 is converted to a sequence of target locations on a cognitive map. Based on the current and target position, a sequence of motion types from the motion pool is generated to guide the robot from its current position to the target position. The navigation combines egocentric local obstacle avoidance and the allocentric global cognitive map. Local navigation is based on data collected directly from sensors, and global navigation is the path planning inside a self-generated global map.

To put the architecture in the real world, a task named “serving milk tea to guest” was performed. We simplified a living-room environment into a 4 m × 4 m maze, as shown in figure 15.8a. The “cup,” “tea,” “hot water,” and “milk” were placed in different locations. The actual experiment with trajectory data is shown in figure 15.8b.

This system offers the capabilities of navigating and mapping in a spatial environment as well as storing and retrieving high-level episodic memories and can be applied to solve high-level service robot tasks. This work would also contribute to developmental robotics by providing a neurophysiological cognitive architecture.

15.4 Conclusion

In this chapter, we presented the development history and the state-of-the-art and elementary components of spatial navigation from the bioinspired perspective, mainly focusing on spatial cells, the cognitive map, and navigation. Next we list some valuable future research directions in biologically inspired spatial cognition and navigation as references for readers.

Multi-map mechanism Rat studies indicate that the brain hosts multiple cognitive maps representing different subsets of the environment at different times and scales. Maps




Figure 15.8

(a) The maze environment for the “serving milk tea to guest” experiment. The required items are placed in different locations in the maze. The arrows indicate the sequence of the action order. (b) The actual experiences trajectory of the mobile robot in the “serving milk tea to guest” task. Task 1 is to get a cup, task 2 is to get the tea, task 3 is to fill the cup with hot water, and task 4 is to add milk.

can be stored and retrieved within a few hundred milliseconds or quickly remapped when environments change or some actions are taken. A future major objective may be to determine how the multiple maps interact with each other and how spatial cells and other factors in the brain contribute to the spatial mapping dynamics. A more comprehensive architecture for space representation and bioinspired navigation will hopefully be developed.

Spatial memory By processing inner and environmental signals, the brain can encode and store spatial information for future retrieval. Research on spatial memory is an ongoing topic in the neuroscience and computer science communities. The input signals can be self-movement signals from the vestibular system, visual information, tactile information, and olfactory and auditory cues. Spatial memory can be stored at several levels, including working memory, short-term memory, and long-term memory. The inner cognitive map in the brain and spatial memory can be integrated to help the robot complete very complex cognitive tasks.

AI and cognitive navigation The firing patterns of spatial neurons in the brain shed new light on spatial navigation research. Whatever form the cognitive map takes, a broad consensus has emerged that spatial cognition and learning can be achieved through a priori and inherent patterns in the brain. In addition, artificial intelligence (AI) studies demonstrate that these a priori patterns can be obtained through pretraining with large data sets and many learning epochs. So the pretraining design may be an important future research direction. Recent studies, such as the curiosity model and the Bert language model, show us how to design pretraining for a priori structure generation.

Nonspatial cognitive task The exploration of spatial cognition provides great inspiration to study many nonspatial cognitive tasks. For example, language reflects a human's ability to use and control signs and can be correlated to spatial cognition: signs correspond to spatial points (or spatial cells), and language corresponds to spatial navigation. If the relationship between spatial cells and navigation can be abstracted as a general cognitive mechanism, maybe we can try to model signs and language from another novel aspect. In an abstracted map with signs, the movement is no longer an action from one point to another in Euclidean space but may be a logical-thinking flow.

Additional Reading and Resources

- This book is the key publication presenting the hippocampal-based approaches to robot navigation and the RatSLAM approach: Milford, Michael. 2008. *Robot Navigation from Nature: Simultaneous Localisation, Mapping, and Path Planning Based on Hippocampal Models*. Vol. 41. Berlin: Springer Science and Business Media.
- This book gives a comprehensive overview of human spatial navigation: Ekstrom, Arne D., Hugo J. Spiers, Véronique D. Bohbot, and R. Shayna Rosenbaum. 2018. *Human Spatial Navigation*. Princeton, NJ: Princeton University Press.
- This paper provides a recent analysis of the neurobiology of mammal navigation: Poulter, Steven, Tom Hartley, and Colin Lever. 2018. "The Neurobiology of Mammalian Navigation." *Current Biology* 28 (17): R1023–R1042.
- Accessible code for RatSLAM: <https://github.com/davidmball/ratslam>.
- Accessible code for NeuroSLAM: <https://github.com/cognav/NeuroSLAM>.

References

- Amit, Daniel J. 1989. *Modeling Brain Function: The World of Attractor Neural Networks*. New York: Cambridge University Press.
- Blair, Hugh T., Adam C. Welday, and Kechen Zhang. 2007. "Scale-Invariant Memory Representations Emerge from Moiré Interference between Grid Fields That Produce Theta Oscillations: A Computational Model." *Journal of Neuroscience* 27 (12): 3211–3229.
- Brooks, Rodney A. 1999. *Cambrian Intelligence: The Early History of the New AI*. Cambridge, MA: MIT Press.
- Brun, Vegard Heimly, Trygve Solstad, Kirsten Brun Kjelstrup, Marianne Fyhn, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. 2008. "Progressive Increase in Grid Scale from Dorsal to Ventral Medial Entorhinal Cortex." *Hippocampus* 18 (12): 1200–1212.
- Burak, Yoram, and Ila R. Fiete. 2009. "Accurate Path Integration in Continuous Attractor Network Models of Grid Cells." *PLoS Computational Biology* 5 (2): e1000291.
- Burgess, Neil. 2008. "Grid Cells and Theta as Oscillatory Interference: Theory and Predictions." *Hippocampus* 18 (12): 1157–1174.
- Burgess, Neil, Caswell Barry, and John O'Keefe. 2007. "An Oscillatory Interference Model of Grid Cell Firing." *Hippocampus* 17 (9): 801–812.
- Buzsáki, György, and Edvard I. Moser. 2013. "Memory, Navigation and Theta Rhythm in the Hippocampal-Entorhinal System." *Nature Neuroscience* 16 (2): 130–138.
- Calton, Jeffrey L., Carol S. Turner, De-Laine M. Cyrenne, Brian R. Lee, and Jeffrey S. Taube. 2008. "Landmark Control and Updating of Self-Movement Cues Are Largely Maintained in Head Direction Cells after Lesions of the Posterior Parietal Cortex." *Behavioral Neuroscience* 122 (4): 827–840.
- Cuperlier, Nicolas, Mathias Quoy, and Philippe Gaussier. 2007. "Neurobiologically Inspired Mobile Robot Navigation and Planning." *Frontiers in Neurobotics* 1:3–3.
- Endo, Yoichiro. 2008. "Anticipatory Robot Control for a Partially Observable Environment Using Episodic Memories." In *2008 IEEE International Conference on Robotics and Automation*, 2852–2859. New York: IEEE.
- Ekstrom, Arne D., Hugo J. Spiers, Véronique D. Bohbot, and R. Shayna Rosenbaum. 2018. *Human Spatial Navigation*. Princeton, NJ: Princeton University Press.
- Erdem, Uğur M., and Michael E. Hasselmo. 2014. "A Biologically Inspired Hierarchical Goal Directed Navigation Model." *Journal of Physiology-Paris* 108 (1): 28–37.
- Fleischer, Jason G., Joseph A. Gally, Gerald M. Edelman, and Jeffrey L. Krichmar. 2007. "Retrospective and Prospective Responses Arising in a Modeled Hippocampus during Maze Navigation by a Brain-Based Device." *Proceedings of the National Academy of Sciences of the United States of America* 104 (9): 3556–3561.
- Fuhs, Mark C., and David S. Touretzky. 2006. "A Spin Glass Model of Path Integration in Rat Medial Entorhinal Cortex." *Journal of Neuroscience* 26 (16): 4266–4276.
- Fyhn, Marianne, Sturla Molden, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. 2004. "Spatial Representation in the Entorhinal Cortex." *Science* 305 (5688): 1258–1264.
- Hafting, Torkel, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. 2005. "Microstructure of a Spatial Map in the Entorhinal Cortex." *Nature* 436 (7052): 801–806.
- Hebb, Donald O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley and Sons.
- Hu, Jun, Miaolong Yuan, Huajin Tang, and Wei Yun Yau. 2016. "Hebbian Learning Analysis of a Grid Cell Based Cognitive Mapping System." *2016 IEEE Congress on Evolutionary Computation*, 1212–1218. New York: IEEE.
- Huang, Weiwei, Huajin Tang, and Bo Tian. 2014. "Vision Enhanced Neuro-Cognitive Structure for Robotic Spatial Cognition." *Neurocomputing* 129:49–58.
- Hubel, David Hunter, and Torsten Nils Wiesel. 1959. "Receptive Fields of Single Neurons in the Cat's Striate Cortex." *Journal of Physiology* 148 (3): 574–591.
- Hubel, David Hunter, and Torsten Nils Wiesel. 1977. "Ferrier Lecture: Functional Architecture of Macaque Monkey Visual Cortex." *Proceedings of the Royal Society B: Biological Sciences* 198 (1130): 1–59.
- Jauffret, Adrien, Nicolas Cuperlier, Philippe Gaussier, and Philippe Tarroux. 2012. "Multimodal Integration of Visual Place Cells and Grid Cells for Navigation Tasks of a Real Robot." *From Animals to Animats 12: 12th International Conference on Simulation of Adaptive Behavior*, 136–145. Berlin: Springer.
- Jockel, Sascha, Daniel Westhoff, and Jianwei Zhang. 2007. "Epirome—a Novel Framework to Investigate High-Level Episodic Robot Memory." In *2007 IEEE International Conference on Robotics and Biomimetics*, 1075–1080. New York: IEEE.

- Kacelnik, Alejandro, and Ian A. Todd. 1992. "Psychological Mechanisms and the Marginal Value Theorem: Effect of Variability in Travel Time on Patch Exploitation." *Animal Behavior* 43 (2): 313–322.
- Killian, Nathaniel J., Michael J. Jutras, and Elizabeth A. Buffalo. 2012. "A Map of Visual Space in the Primate Entorhinal Cortex." *Nature* 491 (7426): 761–764.
- Knierim, James J., and Kechen Zhang. 2012. "Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System." *Annual Review of Neuroscience* 35 (1): 267–285.
- Krichmar, Jeffrey L., Douglas A. Nitz, Joseph A. Gally, and Gerald M. Edelman. 2005. "Characterizing Functional Hippocampal Pathways in a Brain-Based Device as It Solves a Spatial Memory Task." *Proceedings of the National Academy of Sciences of the United States of America* 102 (6): 2111–2116.
- Lever, Colin, Stephen Burton, Ali Jeewajee, John O'Keefe, and Neil Burgess. 2009. "Boundary Vector Cells in the Subiculum of the Hippocampal Formation." *Journal of Neuroscience* 29 (31): 9771–9777.
- McNaughton, Bruce L., Carol A. Barnes, Jason L. Gerrard, Katalin Gothard, Min W. Jung, James J. Knierim, H. Kudrimoti, et al. 1996. "Deciphering the Hippocampal Polyglot: The Hippocampus as a Path Integration System." *Journal of Experimental Biology* 199 (1): 173–185.
- McNaughton, Bruce L., Francesco P. Battaglia, Ole Jensen, Edvard I. Moser, and May-Britt Moser. 2006. "Path Integration and the Neural Basis of the 'Cognitive Map.'" *Nature Reviews Neuroscience* 7 (8): 663–678.
- Méndez, L. A. Torres, and R. Cervantes Jacobo. 2012. "Learning Cognitive Human Navigation Behaviors for Indoor Mobile Robot Navigation." In *COGNITIVE 2012: The Fourth International Conference on Advanced Cognitive Technologies and Applications*, 37–45.
- Milford, Michael. 2008. *Robot Navigation from Nature: Simultaneous Localisation, Mapping, and Path Planning Based on Hippocampal Models*. Vol. 41. Berlin: Springer Science and Business Media.
- Milford, Michael J., Janet Wiles, and Gordon F. Wyeth. 2010. "Solving Navigational Uncertainty Using Grid Cells on Robots." *PLoS Computational Biology* 6 (11): e1000995.
- Milford, Michael J., and Gordon F. Wyeth. 2008. "Mapping a Suburb with a Single Camera Using a Biologically Inspired SLAM System." *IEEE Transactions on Robotics* 24 (5): 1038–1053.
- Milford, Michael, and Gordon F. Wyeth. 2010. "Persistent Navigation and Mapping using a Biologically Inspired SLAM System." *International Journal of Robotics Research* 29 (9): 1131–1153.
- Milford, Michael J., Gordon F. Wyeth, and David Prasser. 2004. "RatSLAM: A Hippocampal Model for Simultaneous Localization and Mapping." In Vol. 1, *Proceedings of the IEEE International Conference on Robotics and Automation*, 403–408. New York: IEEE.
- Monaco, Joseph D., and Larry F. Abbott. 2011. "Modular Realignment of Entorhinal Grid Cell Activity as a Basis for Hippocampal Remapping." *Journal of Neuroscience* 31 (25): 9414–9425.
- Moser, Edvard I., Emilio Kropff, and May-Britt Moser. 2008. "Place Cells, Grid Cells, and the Brain's Spatial Representation System." *Annual Review of Neuroscience* 31 (1): 69–89.
- Moser, Edvard, and May-Britt Moser. 2007. "Grid Cells." *Scholarpedia* 2 (7): 3394.
- O'Keefe, John, and Neil Burgess. 2005. "Dual Phase and Rate Coding in Hippocampal Place Cells: Theoretical Significance and Relationship to Entorhinal Grid Cells." *Hippocampus* 15 (7): 853–866.
- O'Keefe, John, and Jonathan Dostrovsky. 1971. "The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat." *Brain Research* 34 (1): 171–175.
- O'Keefe, John, and Lynn Nadel. 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Poulter, Steven, Tom Hartley, and Colin Lever. 2018. "The Neurobiology of Mammalian Navigation." *Current Biology* 28 (17): r1023–r1042.
- Ranck, J. B. 1984. "Head-Direction Cells in the Deep Cell Layers of Dorsal Presubiculum in Freely Moving Rats." *Society for Neuroscience Abstracts* 10:599.
- Saiki, Luis Yoichi Morales, Satoru Satake, Takayuki Kanda, and Norihiro Hagita. 2011. "Modeling Environments from a Route Perspective." *Proceedings of the 6th International Conference on Human-Robot Interaction*, 441–448. ACM.
- Sargolini, Francesca, Marianne Fyhn, Torkel Hafting, Bruce Mcnaughton, Menno Witter, May-Britt Moser, and Edvard Moser. 2006. "Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex." *Science* 312 (5774): 758–762.
- Savelli, Francesco, D. Yoganasimha, and James J. Knierim. 2008. "Influence of Boundary Removal on the Spatial Representations of the Medial Entorhinal Cortex." *Hippocampus* 18 (12): 1270–1282.
- Sherman, Brynn E., Kathryn N. Graves, and Nicholas B. Turk-Browne. 2020. "The Prevalence and Importance of Statistical Learning in Human Cognition and Behavior." *Current Opinion in Behavioral Sciences* 32:15–20.
- Shim, Vui Ann, Chris Stephen, Naveen Ranjit, Bo Tian, and Huajin Tang. 2013. "A Simplified Cerebellum-Based Model for Motor Control in Brain Based Devices." In *International Conference on Neural Information Processing*. 520–527. Berlin: Springer.

- Shim, Vui Ann, Bo Tian, Miaolong Yuan, Huajin Tang, and Haizhou Li. 2014. "Direction-Driven Navigation Using Cognitive Map for Mobile Robots." In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2639–2646. New York: IEEE.
- Silveira, Luan, Felipe Guth, Paulo Drews-Jr, Pedro Ballester, Matheus Machado, F. Moraes, Nelson Duarte Filho, and Silvia Botelho. 2015. "An Open-Source Bio-inspired Solution to Underwater SLAM." *IFAC-PapersOnLine* 48 (2): 212–217.
- Skaggs, William E., James J. Knierim, Hemant S. Kudrimoti, and Bruce L. McNaughton. 1994. "A Model of the Neural Basis of the Rat's Sense of Direction." *Advances in Neural Information Processing Systems* 7:173–180.
- Solstad, Trygve, Charlotte N. Boccara, Emilio Kropff, May-Britt Moser, and Edvard I. Moser. 2008. "Representation of Geometric Borders in the Entorhinal Cortex." *Science* 322 (5909): 1865–1868.
- Solstad, Trygve, Edvard I. Moser, and Gaute T. Einevoll. 2006. "From Grid Cells to Place Cells: A Mathematical Model." *Hippocampus* 16 (12): 1026–1031.
- Stachowicz, Dennis, and Geert-Jan M. Kruijff. 2012. "Episodic-Like Memory for Cognitive Robots." *IEEE Transactions on Autonomous Mental Development* 4 (1): 1–16.
- Steckel, Jan, and Herbert Peremans. 2013. "BatSLAM: Simultaneous Localization and Mapping Using Biomimetic Sonar." *PLoS One* 8 (1): e54076.
- Tan, Chin Hiong, Huajin Tang, Kay Chen Tan, and Miaolong Yuan. 2013. "A Hippocampus CA3 Spiking Neural Network Model for Storage and Retrieval of Sequential Memory." In *2013 IEEE Conference on Cybernetics and Intelligent Systems*, 134–139. New York: IEEE.
- Tang, Huajin, Weiwei Huang, Aditya Narayanamoorthy, and Rui Yan. 2017. "Cognitive Memory and Mapping in a Brain-Like System for Robotic Navigation." *Neural Networks* 87:27–37.
- Tang, Huajin, Rui Yan, and Kay Chen Tan. 2018. "Cognitive Navigation by Neuro-Inspired Localization, Mapping, and Episodic Memory." *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 751–761.
- Taube, Jeffrey S. 2007. "The Head Direction Signal: Origins and Sensory-Motor Integration." *Annual Review of Neuroscience* 30 (1): 181–207.
- Taube, Jeffrey S., Robert U. Muller, and James B. Ranck. 1990. "Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. I. Description and Quantitative Analysis." *Journal of Neuroscience* 10 (2): 420–435.
- Thrun, Sebastian. 2002. "Probabilistic Robotics." *Communications of the ACM* 45 (3): 52–57.
- Tian, Bo, Vui Ann Shim, Miaolong Yuan, Chithra Srinivasan, Huajin Tang, and Haizhou Li. 2013. "RGB-D Based Cognitive Map Building and Navigation." In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1562–1567. New York: IEEE.
- Tolman, Edward C. 1948. "Cognitive Maps in Rats and Men." *Psychological Review* 55 (4): 189–208.
- Trappenberg, Thomas P. 2002. *Fundamentals of Computational Neuroscience*. Oxford: Oxford University Press.
- Tulving, Endel, and Hans J. Markowitsch. 1998. "Episodic and Declarative Memory: Role of the Hippocampus." *Hippocampus* 8 (3): 198–204.
- Wills, Tom J., Colin Lever, Francesca Cacucci, Neil Burgess, and John O'Keefe. 2005. "Attractor Dynamics in the Hippocampal Representation of the Local Environment." *Science* 308 (5723): 873–876.
- Yartsev, Michael M., Menno P. Witter, and Nachum Ulanovsky. 2011. "Grid Cells without Theta Oscillations in the Entorhinal Cortex of Bats." *Nature* 479 (7371): 103–107.
- Yoshida, Motoharu, and Michael E. Hasselmo. 2009. "Persistent Firing Supported by an Intrinsic Cellular Mechanism in a Component of the Head Direction System." *Journal of Neuroscience* 29 (15): 4945–4952.
- Yu, Fangwen, Jianga Shang, Youjian Hu, and Michael Milford. 2019. "NeuroSLAM: A Brain-Inspired SLAM System for 3D Environments." *Biological Cybernetics* 113:515–545.
- Yuan, Miaolong, Bo Tian, Vui Ann Shim, Huajin Tang, and Haizhou Li. 2015. "An Entorhinal-Hippocampal Model for Simultaneous Cognitive Map Building." In *AAAI'15 Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 586–592. Menlo Park, CA: AAAI Press.
- Zeng, Taiping, and Bailu Si. 2017. "Cognitive Mapping Based on Conjunctive Representations of Space and Movement." *Frontiers in Neurobotics* 11:61.
- Zeng, Taiping, Fengzhen Tang, Daxiong Ji, and Bailu Si. 2020. "NeuroBayesSLAM: Neurobiologically Inspired Bayesian Integration of Multisensory Information for Robot Navigation." *Neural Networks* 126:21–35.
- Zilli, Eric A., and Michael E. Hasselmo. 2010. "Coupled Noisy Spiking Neurons as Velocity-Controlled Oscillators in a Model of Grid Cell Spatial Firing." *Journal of Neuroscience* 30 (41): 13850–13860.

16 Cognitive Robot Manipulation

Yiming Jiang and Chenguang Yang

16.1 Introduction

Manipulation is commonplace among animals and humans. Human beings can effectively manipulate objects with different shapes, weights, sizes, and materials in a variety of tasks such as writing, carrying, pushing, throwing, and rolling. The ability to grasp and manipulate objects is one of the most fundamental human skills. However, these might not be that easy for robots. Robotic manipulation represents the manner in which robots interact with objects, like reaching and grasping an object, picking and placing, opening a door, folding laundry, and so forth (Billard and Kragic 2019). The manipulations are subject to the laws of physics since every manipulation involves a physical robot-environment interaction. The robot must first enter a state in which it can change the state of the object and then apply the desired manipulation. These requirements can be expressed as a set of nonholonomic constraints that define how the robot moves through the entire state space based on its different interactions with the environment. In the past decades, robotic scientists have done extensive research on robot manipulation (see Murray et al. 1994; Siciliano and Khatib 2016; Asada and Slotine 1986; Yang, Luo, et al. 2018; Jiang et al. 2017).

Although a huge effort has been made to advance robotic mechanisms, perception, and control, robot manipulation is far inferior to human in terms of dexterity. To date, robots are still unable to manipulate deformable objects or carry out a screwing task with ease. Therefore, people are seeking more intelligent methods for robot manipulation. To improve the skills of robot manipulation, one natural idea is to understand human manipulation skills and then transfer them to robots. With the prospect of enabling robots to manipulate with humanlike dexterity in scenarios such as sorting, picking and placing objects, folding laundry, and performing house chores, research on the transfer of human skills to robots has attracted considerable attention in the community of roboticists (Yang, Zeng, et al. 2017). This requires robots to learn from human movements and to perform motion planning and to be controlled in a humanlike way to accomplish these actions with dexterity (Tsarouchi et al. 2016). The transferring of human skills to robots can be achieved through 1) learning and 2) perception/cognition of actions of humans. To capture the data from humans, several methods are available: 1) body sensors and 2) visual perception. There

would be many advantages to this human-robot dynamic transfer: safety, compliant interaction with humans, and an environment with low contact force, fewer trajectory errors, and less time spent on robot training.

The objective of this chapter is to introduce the recent state-of-the-art cognitive robot manipulation using advanced sensors to learn from human behavior. Studies of human motor behavior have shown that the central neural system (CNS) can adapt force and impedance in order to interact with the environment optimally. Muscle activities regulated by the CNS can be represented by surface electromyography (sEMG) measured by electrodes attached to the skin (Hermens et al. 2000). They reflect human muscle activation, which represents human joint motion, force, stiffness, and so on and has been employed in robot manipulation tasks (Osu et al. 2002; Ray and Guha 1983). A brain-computer interface (BCI) system is another pathway to connect the human to robots. A BCI system collects a subject's electroencephalograph (EEG) signals, analyzes them, and classifies them to indicate the subject's intention. To help people better understand the function of the brain, BCI can be used to communicate with and control external robotic systems using mental activity (Wolpaw et al. 2000). The EEG signals record the electrical activity of the brain, which can reflect the cortical electrical activity (Guler and Ubeyli 2007). The first report of human EEG was published in 1929 by Hans Berger (Collura 1993), and since then scientists and psychologists have produced a great deal of knowledge regarding EEG, especially in the neuroscience area. Recently, BCI technologies have developed rapidly. They have a wide range of applications in the field of control interfaces (Kosmyna et al. 2016), patient rehabilitation (Young et al. 2014), entertainment (Folgeri and Zampolini 2015), brain cognition, and more. BCI provides a direct pathway to connect the human brain with external devices. This advantage makes it appropriate to combine with a robot system. In Zhao et al. (2015), steady-state visual evoked potentials (SSVEP) BCI was employed to control a humanoid robot. In Geng, Gan, and Hu (2010), a self-paced online BCI was developed for a mobile robot. In Tsui, Gan, and Hu (2011), a motor imagery BCI was designed to control a wheelchair.

Recently, rapid developments in artificial intelligence and deep learning have provided powerful tools for robotic cognitive manipulation. Deep learning has empowered robots to learn various skills, such as pushing (Yuan et al. 2019), grasping (Jang et al. 2018; Kalashnikov et al. 2018), inserting (Lee et al. 2019), and manipulating deformable objects (Matas, James, and Davison 2018), and plays an important role in the strategic planning of subsequent actions. Thus, deep-learning methods have been widely used in robotics. Deep reinforcement learning has been popular since its study in the game of Go (Silver et al. 2016) and in video games (Mnih et al. 2015). Reinforcement-learning algorithms can be grouped into two categories depending on whether the action is continuous (Lillicrap et al. 2015) or discrete (Mnih et al. 2015). As the algorithms continue to improve, there are a growing number of research results on the application of algorithms in robotics. Yuan et al. (2019) examined nonprehensile rearrangement based on deep Q-learning, pushing an object to the predefined goal pose in an environment with obstacles. Nair et al. (2018) utilized a variational autoencoder to encode the input image, calculate the reward based on the Euclidean distance of the encoded vector, and verify this algorithm in an experiment with reaching and pushing. Liang, Lou, and Choi (2019) studied how to attach a flat object to a wall to enable a grasp from the side.

For grasping detection and classification, Redmon et al. (2016) proposed YOLO (you only look once) to classify objects with high accuracy and localize the recognized objects with coordinates of the bounding box simultaneously in real time. Inspired by simultaneous implementation, Trottier, Giguere, and Chaib-Draa (2017) used a residual convolutional neural network to predict the confidence map and the rectangle for grasping single objects in an image. However, for all-purpose utility, a robot should be able to grasp objects in cluttered scenes. A generative grasping convolutional neural network (GG-CNN) was proposed by Morrison, Corke, and Leitner (2018) and used the depth image to predict grasp quality and grasp the pose of every pixel. Zhang, Lan, et al. (2018) developed a region of interest (ROI)-based detection system that can grasp objects in a pile-of-objects scene. All their work and more (Chu and Vela 2018; Kumra and Kanan 2017; Redmon and Angelova 2015) made full use of convolutional neural networks, which needed no preprocessing and could automatically extract features.

In order to make robots manipulate similarly to humans, robots also need to be taught to learn specific skills (Yang, Zeng, et al. 2018; Yang, Zeng, et al. 2019; Yang, Chen, He, et al. 2019; Yang, Chen, Wang, et al. 2019). Traditional robot-teaching methods require professionals to use a teaching pendant to program (Billard et al. 2008). The tasks are programmed with one or a set of discrete movements—that is, point-to-point motions. For example, a grasp-and-place task can be regarded as a combination of several discrete movements: 1) moving the gripper to the object, 2) grasping the object, and 3) moving the gripper to the target position. However, this approach is time-consuming and inefficient to adapt to the work, which requires frequently updated skills. Compared with traditional methods, teaching by demonstration (TbD) or programming by demonstration (PbD) is an efficient way to reduce the complexity of enabling the robot to perform new tasks (Billard and Calinon 2008; Schaal 1999). In a PbD task, a human tutor demonstrates a task, and then a robot learns the motions. The correspondence problem is evident in how the robot imitates the human tutor (Dautenhahn and Nehaniv 2002) movement. One of the solutions is guiding the robot by hand or teleoperating the robot with motion sensors—for example, the motion sensor Kinect, produced by Microsoft. Through human motion capture, visual techniques and devices can be used to enhance the performance of teleoperation-based TbD (Peng et al. 2016). Visual teleoperation techniques, therefore, allow the robots to be programmed directly by learning humanlike manipulation skills from a skillful demonstrator, also known as teleoperation-based TbD.

On the other hand, once human motion data are recorded, we need to build a motion model to transfer them to robot manipulation. The dynamic systems (DS) method, which uses a first-order dynamic system to encode trajectories, has been widely used in motion modeling (Mülling et al. 2013). Hu et al. (2015) proposed a method to learn stable motions from human demonstrations. The implicit mapping of the dynamic system is learned through training a neural network model, the extreme learning machine (ELM). Global stability is ensured by imposing constraints derived from a Lyapunov function on the ELM. This method shows good performance in convergence and generalization. A stable estimator of dynamical systems (SEDS) is another approach of PbD with dynamic system representation (Khansari-Zadeh and Billard 2011). The stability at the target is confirmed by the theory of Lyapunov stability. The difference of the DS method is that the unknown function is modeled as a Gaussian mixture model (GMM) in order to encode the joint

distribution. The GMM makes it possible to account for the features of many demonstrations of a task. As mentioned in Calinon et al. (2012), the GMM combined with Gaussian mixture regression can provide additional information when learning from multiple demonstrations. The probabilistic framework of statistical learning is a powerful tool in PbD (Calinon 2018; Rozo et al. 2013). The dynamic movement primitive (DMP) is widely applied to motion modeling. A DMP represents the movement trajectory by using a spring-damper system coupled with a nonlinear term (Schaal, Mohajerian, and Ijspeert 2007). The inherent stability of the spring-damper system ensures that the generated motion is stable enough to reach the target and robust to perturbation. Therefore, it is unnecessary to impose extra constraints on the model. The generalization ability of a DMP is acquired from a single demonstration. This provides sufficient room to improve the generalization ability.

Next, we will introduce some state-of-the-art techniques for acquiring physiological signals and human motion behavior and their applications, such as EEG visual-system-based robot object picking, visual teleoperation, and motion sensors based on human skill learning and generalization. Moreover, techniques of robot learning by demonstration and skill generalization approaches in robot manipulation are detailed. Additionally, some recent progress in deep learning and machine learning for robot manipulation is introduced. Finally, a brief conclusion will be presented to summarize these works.

16.2 Techniques to Capture Human Information on Robot Manipulation

16.2.1 Surface Electromyography Signals

Ideally, to be used for robot manipulation, sEMG signals reflect human muscle activation and embed rich information about human joint motion, force, stiffness, and so on. Generally, sEMG signals can be processed into two divisions: finite class recognition serials and continuous control reference. The former usually refers to pattern recognition, such as hand posture recognition (Chu et al. 2006; Khezri and Jahed 2007), and such data serials are usually used for switch control, while the latter refers to extracting continuous force, stiffness, and even motion serials from sEMG signals, which reflect the variations of human limb kinematics and dynamics during limb movement or pose maintenance. Furthermore, the relationship between sEMG and stiffness, force, and motion is approximately linear, and thus biocontroller design tends to be simple in sEMG-based robot control systems.

16.2.2 Electroencephalograph Signals

The EEG application focuses on two types of signals: evoked potential (EP) and spontaneous signal modulation. Evoked potential, including visually evoked potential (Müller-Putz et al. 2005) and P300 event-related potential (Rebsamen et al. 2007), is the electrical activity of the nervous system, stimulated by internal and external stimuli. An EEG signal acquisition device is shown in figure 16.1.

When a certain area of the cerebral cortex is activated, metabolism and information processing in this region will increase, leading to the amplitude reduction or blocking of the brain waves, especially in the alpha and beta rhythm. This electrophysiological phenomenon is called event-related desynchronization (ERD). On the contrary, when this



Figure 16.1

The EEG signal acquisition device Neuroscan. EEG raw data are collected by a Neuroscan device with twenty-seven electrode channels (*left*). The EEG sampling rate is set to 250 Hz, and a band-pass filter of 0.5–40 Hz is used in the SCAN 4.5. *Source:* From Wu et al. 2017.

region is at rest, the brain waves will show an obvious increase in amplitude, which is called event-related synchronization (ERS). Studies have shown that in unilateral limb movement or motor imagery, the contralateral side of the brain produces ERD, while the ipsilateral side of the brain produces ERS. It means that if we image right hand movement, the power of the EEG signals will be reduced on the left side of the brain, increased on the right side of the brain, and vice versa. According to the ERS/ERD phenomenon, we can classify the EEG signals into two categories, imaging left and right hand movement. The BCI system can extract the thinking activity information inside the brain using specific measurement technology and then analyze the real intent of the human brain contained in the information through the embedded platform and convert it into the control command to the external device to realize the goal of the human brain directly controlling the peripheral device. The operator can use this to operate the mechanical arm as needed, which overcomes the limitation of using traditional physical buttons as control ports.

16.2.3 Visual Sensor

Visual sensors are widely employed in robot manipulation and control in applications such as visual servo control and object detection. In this part we briefly introduce several kinds of visual sensors and will be explaining their specific usage scenarios later. The representative visual sensors include the Bumblebee2, ZED, and Kinect sensors. The Bumblebee2 is a stereo camera with two CCD strictly paralleled cameras. At one time, Bumblebee2

captures two photos of the robot and the other objects with its two sensors, respectively (Yang, Ma, and Fu 2016). The ZED stereo camera is a passive-depth camera that consists of two RGB cameras with fixed alignment (Yang, Ma, and Fu 2016). The Kinect sensor is widely applied to full-body three-dimensional (3D) motion capture and facial recognition and more (Xu et al. 2017; Wu et al. 2012).

A visual servo system is used in the study of robotic manipulator control by using visual sensors to reconstruct realistic scenes and objective detection. Bumblebee2 can obtain the depth information and a 3D model of the scene in real time. ZED has a high frame rate, a wide field of view, and the ability to run in multiple environments. Kinect has real-time motion capture, image recognition, microphone input, speech recognition, and other functions. The choice of vision sensor will be explained later before the specific use.

16.3 Robot Manipulation by Skill Transferring

16.3.1 Robot Manipulation Using EEG Signals

In this part we introduce an innovative robot arm control method of 3D space manipulation using EEG signals. The realized system is designed to allow users to teleport robots to perform tasks using EEG signals with no hands or feet involved. This system contains three parts: the BCI, visual feedback, and the robot control platform.

System overview

The working mechanism of the system is depicted in figure 16.2. We employ a Bumblebee2 to detect the 3D coordinates of the target objects and the end effector of the robot in real time. These coordinates will be used to display the representations of the objects and the robot in a two-dimensional (2D) plane on the screen. The user looks at the screen, decides the direction of the robot, and generates the specific EEG signals immediately. These signals are collected by Neuroscan and analyzed by a server computer to be converted into robot commands. Then the end effector of the robot will move accordingly in 3D space. As the end effector moves, we can, in the meantime, get position feedback information from the screen to decide the motion of the next movement. In the experiment, Bumblebee2 keeps taking photos with its left and right cameras, from which we construct a disparity map to get one object's depth information. Then we can detect its position in the image and read the object's depth information from the disparity map.

To represent the position of the target objects and the end effector of the robot in a 2D plane on the screen, we should decompose 3D coordinates into several lower-dimension coordinate systems. The representation rectangles' positions are based on the x and y coordinates of the objects and robot hand, which are captured by Bumblebee2 in real time.

Object picking using BCI

To control the robot hand to manipulate in 3D space, six commands are needed: up, down, forward, backward, left, and right. However, our BCI system only offers two kinds. In order to employ the BCI system, coordinate decomposition for the robot control system is needed. At first, the 3D coordinate system will be decomposed into a 2D plane and a z -axis where the 2D plane is parallel with the desktop, and the z -axis indicates the vertical

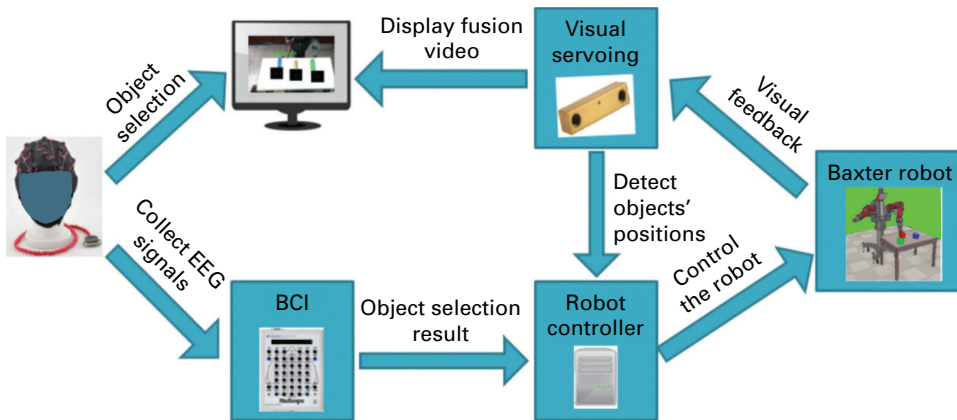


Figure 16.2
Robot control system overview. *Source:* From Yang, Wu, et al. 2018.

direction. As the height of the desktop is changeless, we do not need to manually adjust the z coordinates of the robot hand in picking. It can be designed as a fixed mode. So we need to adjust the position of the robot hand in the xy plane. We continue to divide this plane into two parts—that is, the x -axis and the y -axis. We first adjust the robot hand in the direction of the x -axis to be coincident with the destination's x -axis. Next, we adjust the y coordinate to be the same as the destination. Then we finish the xy plane's adjustment, and the x and y coordinates of the robot's end effector are the same as the destination.

Employing the method of coordinates decomposition mentioned above, the main process of object picking is shown in figure 16.3a. In the first step, the robot's end effector stays on a 2D plane above and parallel to the desk. Bumblebee2 captures the positions of both the robot end effector and the objects and then abstracts and displays them on the screen. Based on this visual feedback, the subject uses their mind to control the robot hand through the BCI system to move to the destination, which is directly above the target object. When the subject finishes this procedure, the control system will detect it, and the end effector will go down to pick up the target object.

The detailed process of how to adjust the robot hand to be directly above the target object is shown in figure 16.3b. As figure 16.3b.I shows, in the first step of the experiments the target objects are abstracted as several rectangles. And the robot's end effector is abstracted as a line. The subject uses their mind to control the end effector of the robot to move along the direction of the x -axis. The robot line on the screen will move toward the left or the right according to the x coordinate of the end effector. As figure 16.3b.II shows, if the robot line reaches one of the rectangles, which means that the x coordinate of the end effector and the target object are the same, the color of the rectangle will change to blue. In this step, if the user keeps on moving the end effector of the robot in the same direction as before, the robot line will go out of the rectangle, and then the system will go back to the first step. Otherwise, if the user changes the direction, the robot line will remain in the rectangle without moving. If the system detects that the robot line has remained in the rectangle for more than two seconds, it will go to the next step.

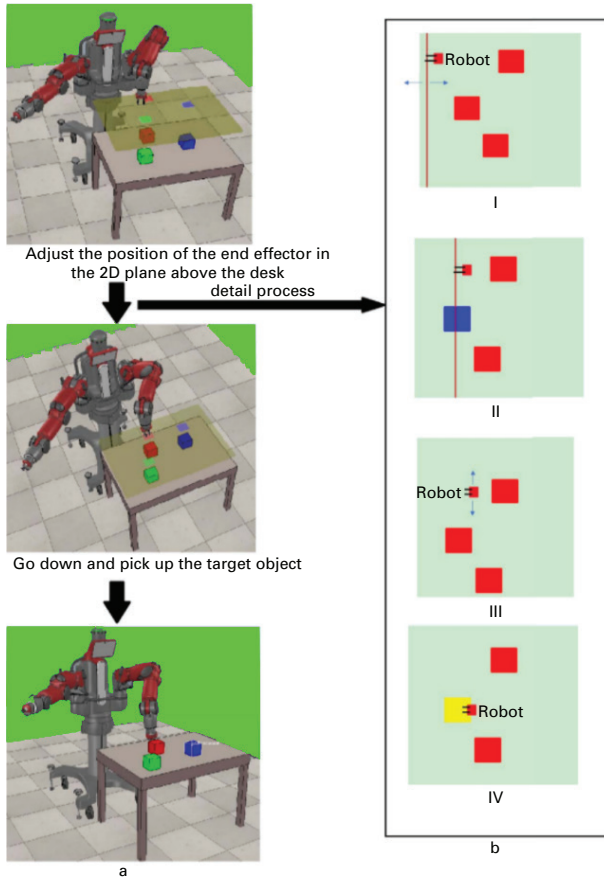


Figure 16.3
Process of picking up an object.

Figure 16.3*b.III* shows the end effector of the robot represented as a little square on the screen in this step. The square is just able to move up and down, which is the direction of the y -axis of the real world. The user controls their mind to move the robot's end effectors along the direction of the coordinate axis. As figure 16.3*b.IV* shows, if the square enters one of the rectangles, the color of the rectangle will change to yellow. Similar to the second step, if the user keeps going in the same direction, the square will leave the rectangle, and the system returns to the third step. Otherwise, if the square remains in the rectangle for more than two seconds, the target object will be selected, and the robot will prepare to pick up this object. As the x coordinate and the y coordinate of the end effector and the target object are the same after the user completes the steps above, the end effectors of the robot just need to move along the direction of the z coordinate and use its clamper to pick up the object.

In these experiments the subject wears an EEG cap, and the machine will collect their EEG data for analysis. Following the steps we introduced before, the subject watches the video on the screen and decides which directions to move the robot block on the screen. When the subject makes their decision, their brain will generate specific EEG signals of

some mode. The BCI system analyzes the signals, classifies them into different commands, and sends them to the robot control platform. The robot's system receives these commands and controls the robot's manipulator to go in the given direction.

In the experiments, the BCI achieves an accuracy rate of about 70 percent. For stable control, we employ a control strategy using the maximum probability principle. We collect a series of BCI commands and analyze them to infer which command may be the real decision of the subject for the highest probability. By employing the mean, we can avoid some errors in the BCI part. The experiment result has shown a high recognition rate and a high efficiency of the robot system. After training several times, three subjects have been able to use the system to control the robot to pick up within one minute.

16.3.2 Robot Manipulation Based on Visual Teleoperation

This section introduces an example of visual teleoperation, which allows the robot manipulation to be controlled by a human demonstrator. Specifically, in the beginning the demonstrator controls the robot by visual interaction. A learning algorithm based on a radial basis function (RBF) network is used to transfer the demonstrator's motions to the robot. Several simulation experiments have been carried out to verify the effectiveness of this advanced method.

The virtual teleoperation system

The virtual teleoperation system, which can simulate the real system, consists of a human demonstrator, a Kinect sensor, and a computer with V-REP, as shown in figure 16.4. Separately, the human should make a demonstration of a specific task; Kinect is applied to capture the human body motion, and V-REP is used to build a virtual work environment and robot.

Space vector approach to calculate human joint angle

Calculating the human joint angle is the key to controlling the Baxter robot by Kinect. Kinect V2 is able to capture twenty-five joint points of a human body in Cartesian space. In a 3D space, the distance between two points $A(x_1, y_1, z_1)$ and $B(x_2, y_2, z_2)$ can be calculated

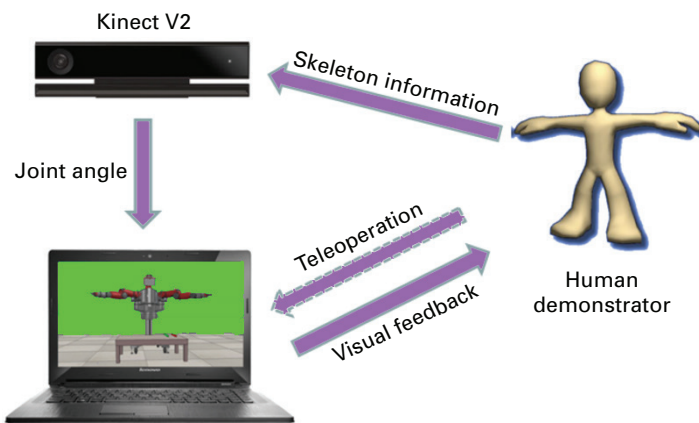


Figure 16.4

The architecture of the teleoperation system. *Source:* From Xu et al. 2017.

by the equation $d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2}$. The vector \overline{AB} can be expressed as $\overline{AB} = (x_2 - x_1, y_2 - y_1, z_2 - z_1)$, $d = |\overline{AB}|$. In 2D or 3D space, we can use the law of cosines to calculate any desired angle between two vectors. And a joint in a Kinect coordinate can be expressed as a vector. If joint 1 is $[-0.987, 0.564, 0.635]$, and joint 2 is \overline{BC} , the angle between the two joints can be calculated as $\cos(\overline{AB}, \overline{BC}) = \frac{\overline{AB} \cdot \overline{BC}}{|\overline{AB}| \cdot |\overline{BC}|}$.

Using the above equations, we can convert all the coordinates detected by Kinect to the corresponding vectors, and the respective angles of the joints in 3D space can be calculated by the law of cosines $\cos(\overline{AB}, \overline{BC})$.

According to the above equation, we obtain all the location coordinates from Kinect, then we build a geometric model of the human left arm as shown in figure 16.5. The directed straight lines OX , OY , and OZ form a coordinate system in the Cartesian space of Kinect. From three points O , E , and F , we can get the vectors \overline{OE} and \overline{EF} . Finally, we will calculate the shoulder pitch angle $\angle OEF$. Using the same method, we can get the elbow pitch $\angle EFG$. And by projecting points D , O , and F to the plane XOZ , we can calculate the shoulder yaw angle $\angle KOJ$. So we solve for the angle of shoulder roll $\angle LEM$. Using the same method, we can calculate the elbow roll, which is the angle between \overline{LE} and \overline{GN} , and the hand yaw, which is the angle between \overline{GN} and \overline{GQ} . To make data processing simple, we control the shoulder joint S_0 , S_1 and the elbow joint E_1 by using a space vector approach.




Figure 16.5
The geometry model of the human left arm.

TbD based on an RBF network

Using the virtual system, we can control Baxter's movement according to a trajectory. And these point-point motions can be modeled as a dynamical system, which can be expressed by a first-order autonomous ordinary differential equation

$$\dot{x} = f(x) + \varepsilon$$

where x denotes the robot's end-effector position or joint angles, and \dot{x} is the first derivative of x . The data set is $\{x, \dot{x}\}$, and ε is a zero mean Gaussian noise. The goal is to obtain an estimation of \hat{f} from f . To achieve this goal, we use a method based on a radial basis function (RBF) network, which has universal approximation and regularization capabilities. If the radial basis function can be suitably chosen, the RBF network will approximate any continuous function arbitrarily (Wu et al. 2012).

The process of training and trajectory reproduction

In the demonstration phase, the robot has completed certain tasks, and the joint angle is recorded as training data. Then the training data are sent to an RBF neural network to get a new set of joint angles. We implement the learning algorithm based on the RBF network in MATLAB. The function *newrb* in MATLAB is mainly used to build the RBF network. The relationship between mean square error (MSE) and the number of hidden neurons are shown in table 16.1. Using the constructed RBF network to approximate the robot trajectory, we obtain three groups of output data.

In the next step, the data is sent to the robot in V-REP by MATLAB. Then the robot will reproduce the trajectory it learned from a human demonstrator. Using the virtual system, we can control Baxter's movement according to this trajectory.

To verify the effectiveness of this TbD method, a simulation scene is designed in the V-REP. The scene consists of a Baxter robot, a desk, and some rectangular building blocks. As shown in figure 16.6, using the Kinect, a human demonstrator controls the Baxter robot as it knocks over a building block, and then the robot's arm returns to its original position. This action will be performed many times. During this, the robot joint angle in this process is recorded at regular intervals.

Because each simulation when controlling the robot is different, the number of data sets recorded in the experiment is also different. We randomly interpolate the experimental data. Using these processed joint angles to control the robot, the robot can reproduce the same trajectory, which confirms that this method has no effect on the effect of the robot's trajectory. Therefore, a sample with the dimension 142×3 is obtained from each simulation after data processing.

Table 16.1

The relationship between MSE and the number of hidden neurons

Neurons	50	100	150	200	250	300
MSE(S0)	0.0122	0.0088	0.0081	0.0079	0.0079	0.0079
MSE(S1)	0.0391	0.0183	0.0117	0.0107	0.0106	0.0106
MSE(E1)	0.0073	0.0072	0.0071	0.0070	0.0070	0.0070




Figure 16.6

Controlled by a human with a Kinect, Baxter can knock over a building block on a desk.

In the next step, the data is sent to V-REP through MATLAB to control the Baxter robot. As shown in figure 16.7, the Baxter in V-REP can autonomously complete the task taught by a human demonstrator.

16.4 Human-to-Robot Skill Transfer and Generalization for Manipulation

In this section we introduce skills transfer and a generalization approach for robot manipulation to learn point-to-point motions from human demonstrations. This work enables the model to learn from a set of demonstrations of a task and extract better motions. Additionally, the generalizability of the original dynamic movement primitive (DMP) is inherited, including the ability of spatial scaling and temporal scaling. We apply the method to the virtual Baxter robot and use the Kinect, a motion sensor, to capture the human demonstrations.

A DMP consists of a spring-damper system and an external forcing term. The model is defined as follows:

$$\begin{aligned}\tau\dot{\omega} &= k(g - \theta) - c\omega + (g - \theta_0)sf(s) \\ \tau\dot{\theta} &= \omega\end{aligned}\quad (16.1)$$

where $\theta \in R$ is the joint angle, $\omega \in R$ is the angular velocity of the joint, $g \in R$ is the goal, $\theta_0 \in R$ is the start angle, $\tau > 0$ is the temporal scaling factor, $k > 0$ is the spring constant, $c > 0$ is the damping coefficient, $f: R \rightarrow R$ is assumed to be a nonlinear continuous bounded function, and $s \in R$ is the state of a first-order dynamic system:

$$\tau\dot{s} = -\alpha_s s \quad (16.2)$$

where $\alpha_s > 0$ is the time constant. This system is referred to as a canonical system. It is introduced to remove the nonlinear function's dependence on time so that the whole system




Figure 16.7

The experiment results: through learning and training via an RBF network, Baxter can autonomously knock over a building block.

(16.1) is autonomous. The state s is regarded as a phase variable. It is monotonically decreasing and will converge to zero. Generally, the initial value is chosen as $s_0 = 1$. More detail about the DMP can be found in Schaal, Mohajerian, and Ijspeert (2007). The generalization ability of the DMP is acquired from a single demonstration, which provides enough space to improve the generalization ability.

With spatial and temporal scaling, we can get a similar motion through modulating the parameters g , θ_0 of the DMP. By setting the factor τ , we can adjust the speed of the generated motion. The issue with the DMP is how to determine the nonlinear function $f(s)$ —that is, the weights γ . The function approximation problem can be solved using locally weighted regression (Atkeson, Andrew, and Schaal 1997). However, this method can only be used for a single demonstration. In order to model multiple demonstrations, the GMMs can be employed (Calinon, Guenter, and Billard 2007).

16.4.1 Learning from Multiple Demonstrations

For the given demonstrations $\{\theta_{t,n}, \dot{\theta}_{t,n}, \ddot{\theta}_{t,n}\}_{t=0,n=1}^{T_n,N}$, where $\theta_{t,n} \in R$ is the joint angle, T_n is the duration of the demonstrations, and N is the number of demonstrations, we first need to calculate the data set from $\{s_t, f_{t,n}\}_{t=0,n=1}^{T_n,N}$ (see figures 16.8a and 16.8b). The $s_t \in R$ is the state of the system at time step t , and the $f_{t,n} \in R$ is calculated through substituting $s_t, \theta_{t,n}, \dot{\theta}_{t,n}, \ddot{\theta}_{t,n}$ into the first differential equation of system (16.1). When $N=1$ —that is, a single demonstration is available—the function $f(s)$ can be learned from the data set $\{S_t, f_{t,1}\}_{t=0}^{T_1}$ using locally weighted regression (LWR). However, this method is not suitable when learning from several demonstrations at the same time. To solve this problem, we use GMMs to model these demonstrations. For convenience, we use $\{s, f\}$ to represent the data set $\{S_t, f_{t,n}\}_{t=0,n=1}^{T_n,N}$ in the remainder of the chapter.

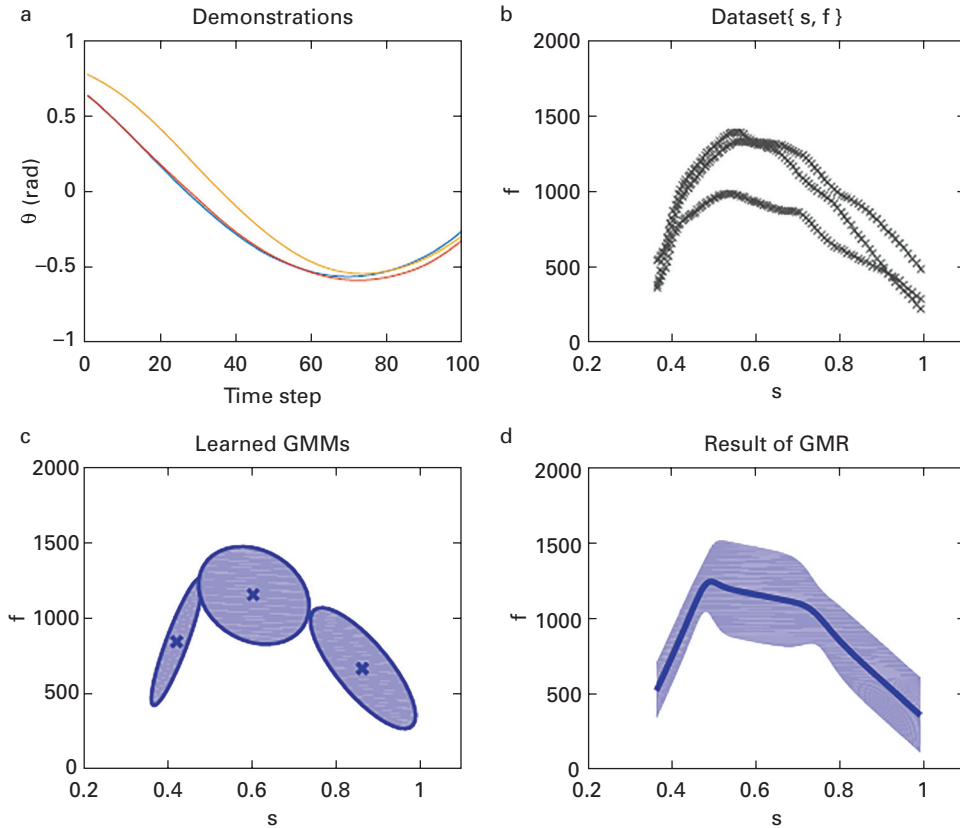


Figure 16.8

The learning process of GMMs/GMR. (a) Three demonstrations in joint space. (b) The data set $\{s, f\}$ calculated from (a). Here we select one hundred points. (c) The learned GMMs, which encode the joint distribution with three Gaussian models. (d) The learning result of GMR. The estimation \hat{f} is retrieved from the GMMs.

A GMM is a statistical method of probability density estimation. Combined with Gaussian mixture regression, it can be used to estimate nonlinear functions. Here, we use GMMs to encode the joint distribution $P(s, f)$, which is defined as follows:

$$P(s, f) = \sum_{k=1}^K \alpha_k \mathcal{N}(s, f; \mu_k, \Sigma_k) \quad (16.3)$$

$$\sum_{k=1}^K \alpha_k = 1 \quad (16.4)$$

$$\mu_k = \begin{bmatrix} [1, 2] \mu_{s,k} \\ \mu_{f,k} \end{bmatrix}, \Sigma_k = \begin{bmatrix} [1, 2] \Sigma_{s,k} & \Sigma_{sf,k} \\ \Sigma_{fs,k} & \Sigma_{f,k} \end{bmatrix} \quad (16.5)$$

$$\mathcal{N}(s, f; \mu_k, \Sigma_k) = \frac{e^{-0.5([s, f]^T - \mu_k)^T \Sigma_k^{-1} ([s, f]^T - \mu_k)}}{2\pi \sqrt{|\Sigma_k|}} \quad (16.6)$$

where k is the number of Gaussian models, $\alpha_k \geq 0$ is a prior probability, $\Sigma_k \in \mathbb{R}^{2 \times 2}$ is the covariance matrix of the k th Gaussian model, $\mathcal{N}(S, f; \mu_k, \Sigma_k)$ is the Gaussian probability distribution.

The $\alpha_k, \mu_k, \Sigma_k$ are unknown parameters of the models. They can be estimated using the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1977), an iterative method of the maximum likelihood estimation. This algorithm is sensitive to the initial value of the parameters. Thus, the k-means method (MacQueen 1967) should be used to initialize the models. Additionally, the number of models influences the error and the smoothness of the estimation. It can be chosen empirically or be estimated through model selection approaches, such as the Bayesian information criterion (BIC).

16.4.2 Experiments

These experiments are based on a virtual robot platform V-REP, as discussed in the previous section. In our experiments, we use the virtual Baxter robot in its model library. We also use vision techniques to capture a human tutor's action and then transmit the demonstration information to the robot. The states of the human tutor's joints are captured by the Kinect, and these shoulder joints and elbow joints are used in the experiments.

In our experiments, the robot is expected to perform a task of pushing boards off of a table. If the Baxter robot needs to do that with its left arm, it has to try to avoid obstacles. The human tutor will show the robot how to perform this task, and the robot learns from the demonstrations. This type of task is common in our daily lives. If you want to grab something placed in a messy environment, you have to avoid other objects before you touch it. In this situation, the shape of the motion trajectory is important for the completion of the task. Here, the DMP model is utilized to learn and further generate human motion skills such as those described above.

16.4.3 Motion Learning and Generation

In the first experiment, the robot learns how to push the board on its left and how to reproduce the learned motion. As shown in figure 16.9, the Baxter robot imitates the human tutor's motions in order to perform the task. The robot raises its left arm over the pillar

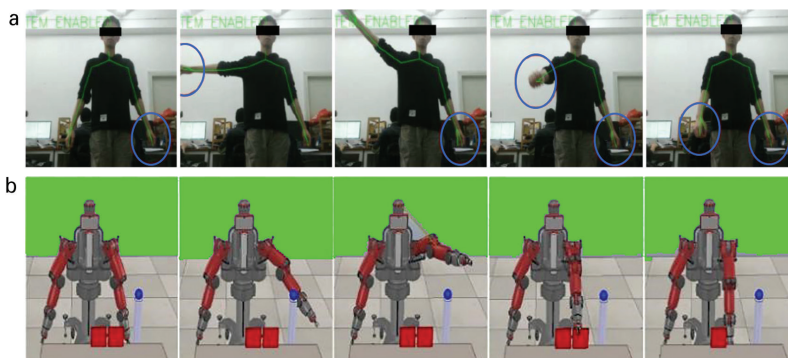


Figure 16.9

(a) A human tutor demonstrates how to push the left board off of the table. (b) The robot imitates the motion of the human tutor.

and then moves its gripper to push the left board off the table. This demonstration is repeated ten times. All of the joint angles that we focus on are recorded throughout the demonstrations. Then the data is used to train the DMP. In order to match the autonomy of the dynamic system, we use a time step to represent the duration of the motion, and all durations of the demonstrations are converted to one hundred steps.

The original DMP is learned from a single demonstration. In order to compare the performance of the two methods, we use those demonstrations to train ten DMPs. We use them to generate new motions without modifying the goal. Then we apply them to the robot to test whether the robot is able to perform the task successfully. In six of the tests, the robot cannot push the left board off of the table or will push both boards off. After modifying the goal of these DMPs to an appropriate position, the robot can complete the task with the regenerated motions.

In the second experiment, we modify the initial angle of one joint to evaluate the stability of the DMP. To evaluate the spatial scaling ability of the DMP, we modify the goal of the motion to another board. The previous goal of three joint angles is $[\theta_1, \theta_2, \theta_3] = [-0.755, 0.652, 0.664]$ (rad). We modify it to $[-0.987, 0.564, 0.635]$ so that the left arm of the robot can reach the right board. We also apply the generated motion to the robot. As figure 16.10a shows, the robot's left arm moves around the pillar and then pushes the right board successfully. Another ability of DMP is temporal scaling. We adjust the temporal scaling factor τ from 1 to 0.5, which can speed up the generated motion. Three joint angles reaching the goals at time step = 50 are shown in figure 16.10b.

16.5 A Brief Introduction to Deep Reinforcement Learning for Robot Grasping and Manipulation

Deep reinforcement learning plays an important role in the strategic planning of sequential actions. Most applications of reinforcement learning in robots are low-level control methods that need long sequences to achieve the goal. The large-scale exploration space and the delayed reward makes it difficult to get training data of high quality, and thus a




Figure 16.10

(a) The motion generated while modifying the goal position to another board. (b) The motion generated while modifying the temporal scaling factor τ from 1 to 0.5.

lot of time is needed to collect data. As presented in Quillen et al. (2018) and Levine et al. (2018), a robot requires more than one hundred thousand grasps to learn the grasping skill without camera calibration. This amount of data requires multiple robots to execute a task over a long period of time, which is costly for most of us, and it is challenging to transfer the skill to different robots. As for deep supervised learning of the grasping skill, Chu and colleagues (e.g., Chu, Xu, and Vela 2018) came up with a proposal based on the faster region-based CNN (RCNN) network to transfer the grasp rectangle detection to object detection, resulting in high classification performance. In Morrison, Corke, and Leitner (2018), they achieved pixel-wise grasp rectangle detection by using a fully convolutional network like Unet to predict the rectangle for every pixel. Without fully connected layers, their network was significantly smaller than other networks. In a task of clearing clustered objects that needs to combine pushing and grasping, we are inspired by an algorithm that maps the image to high-level action instead of continuous action of a low level based on the mapping relation between the image and the workspace (Berscheid, Meißner, and Kröger 2019). One-to-one correspondence between discrete actions and pixels has the ability to make precise decisions but leads to a large network and a long reasoning time.

In Zeng et al. (2018), pushing and grasping were both learned based on reinforcement learning. They used Q-learning to choose discrete actions on pixel-wise and map the pixel coordinates to the real-world location. It should be noted that the more complex the task, the more time is needed for the real robot to collect data by interacting with the environment and for neural network fitting. Especially for grasping, few positive samples and diverse objects mean that hundreds of hours of collecting data is inescapable. Although sim-to-real techniques can ease this problem to some extent, learning to grasp with reinforcement learning is still time-consuming and costly. As for pushing, there are multiple solutions to separate objects. This kind of problem is hard to define manually and doesn't require a very precise solution; hence, it is suitable for reinforcement learning.

In our recent work (Chen, Yang, and Feng 2020), we found that the grasp detection algorithm based on supervised learning was mostly trained on the Cornell Grasping Dataset or the Jacquard Dataset, whose depth image is strikingly different from the depth image in simulation because of different shooting angles. Therefore, we utilized a traditional morphological method in Zhang, Yang, et al. (2018), which can be easily transferred to a virtual image with little change. Compared with their work, the framework applies a policy that outputs continuous action to avoid large action dimensions, and the accurate grasp point detection ensures a high grasp rate of graspable objects. Therefore, the framework is simple in structure but competent for the clutter-clearing task. We have combined pushing based on reinforcement learning and grasping based on a traditional, rule-based grasp-detection algorithm (Zhang, Yang, et al. 2018). We employ the twin delayed deep deterministic policy gradient (Fujimoto, van Hoof, and Meger 2018) to train the policy that determines where to start pushing and the pushing direction according to the current image depth. The pushing direction within 360° is divided into two sides, and we introduce a variable to decide which direction one needs to push toward. The grasp detection is processed with a rule-based method mainly based on the recognition of a minimum bounding convex hull and a minimum bounding rectangle of connected regions. The grasp detection algorithm determines whether an object is graspable and computes the grasp center and the grasp orientation. When performing the task, the pushing

action is executed only when no object is graspable, which is determined by the grasp detection algorithm.

16.6 Conclusion

Robots with arm manipulators are shifting from industrial factories to factors in people's daily lives, such as home services and medical care. One possible solution to achieve these goals is to enable the robots to learn manipulation skills from human behavior. This chapter investigated a number of effective methods to transfer human adaptation skills to robots within a variety of sensor data—for example, physiological signals such as EEG, body motion signals, visual signals, and so on. The PbD approach was proposed to improve the efficiency of a robot learning human motion skills through imitating a human tutor. We combined the DMP with GMMs to enable a robot to learn from a set of demonstrations that are captured by a motion sensor to provide friendly human-robot interaction. Moreover, we introduced some recent progress in deep learning applied to robot manipulation. Future work includes teaching a robot in a more intelligent manner so that a robot can learn more dexterous skills from a human using position, stiffness, and force information.

Acknowledgments

We would like to acknowledge the enormous contributions from Kunlin Guo, Yang Liu, and Chengzhi Zhu during the preparation of this chapter.

Additional Reading and Resources

- A classical book introducing the basic theory and mathematical foundation of robot manipulation: Murray, R. M., Z. Li, S. S. Sastry, and S. S. Sastry. 1994. *A Mathematical Introduction to Robotic Manipulation*. Boca Raton: CRC Press.
- A systematic handbook of robotics, with specific sections on robotic manipulation: Siciliano, Bruno, and Oussama Khatib, eds. 2016. *Springer Handbook of Robotics*. Berlin: Springer.
- A recent and comprehensive book introducing advanced technologies of robotic manipulation, with specific sections on bioinspired robotic manipulation and visual servoing control: Yang, C., H. Ma, and M. Fu. 2016. *Advanced Technologies in Modern Robotic Applications*. Singapore: Springer.
- Physiological signals enhanced manipulation: <https://www.youtube.com/watch?v=rvHluEVSyZw>.
- Skills transfer from human to robot using sEMG: <https://www.youtube.com/watch?v=CCKy88QTkGY>.

References

- Asada, Haruhiko, and J-JE. Slotine. 1986. *Robot Analysis and Control*. Hoboken, NJ: John Wiley and Sons.
- Atkeson, Christopher G., Andrew W. Moore, and Stefan Schaal. 1997. "Locally Weighted Learning for Control." *Artificial Intelligence Review* 11:75–113.

- Berscheid, Lars, Pascal Meißner, and Torsten Kröger. 2019. "Robot Learning of Shifting Objects for Grasping in Cluttered Environments." ArXiv preprint: 1907.11035.
- Billard, Aude, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. 2008. "Robot Programming by Demonstration." In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1371–1394. Berlin: Springer.
- Billard, Aude, and Danica Kragic. 2019. "Trends and Challenges in Robot Manipulation." *Science* 364 (6446): eaat8414.
- Calinon, Sylvain. 2018. "Robot Learning with Task-Parameterized Generative Models." In *Robotics Research*, 111–126. Cham, Switzerland: Springer.
- Calinon, Sylvain, Florent Guenter, and Aude Billard. 2007. "On Learning, Representing, and Generalizing a Task in a Humanoid Robot." *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37 (2): 286–298.
- Calinon, Sylvain, Zhibin Li, Tohid Alizadeh, Nikos G. Tsagarakis, and Darwin G. Caldwell. 2012. "Statistical Dynamical Systems for Skills Acquisition in Humanoids." In *2012 12th IEEE-RAS International Conference on Humanoid Robots*, 323–329. New York: IEEE.
- Chen, Yiwen, Chenguang Yang, and Ying Feng. 2020. "Reinforcement Learning on Robot with Variational Auto-Encoder." In *Proceedings of the 11th International Conference on Modelling, Identification and Control*, 675–684. Springer, Singapore.
- Chu, Fu-Jen, and Patricio A. Vela. 2018. "Deep Grasp: Detection and Localization of Grasps with Deep Neural Networks." ArXiv preprint: 1802.00520.
- Chu, Fu-Jen, Ruinian Xu, and Patricio A. Vela. 2018. "Real-World Multiobject, Multigrasp Detection." *IEEE Robotics and Automation Letters* 3 (4): 3355–3362.
- Chu, Jun-Uk, Inhyuk Moon, and Mu-Seong Mun. 2006. "A Real-Time EMG Pattern Recognition System Based on Linear-Nonlinear Feature Projection for a Multifunction Myoelectric Hand." *IEEE Transactions on Biomedical Engineering* 53 (11): 2232–2239.
- Collura, Thomas F. 1993. "History and Evolution of Electroencephalographic Instruments and Techniques." *Journal of Clinical Neurophysiology* 10 (4): 476–504.
- Dautenhahn, Kerstin, and Chrystopher L. Nehaniv. 2002. "The Agent-Based Perspective on Imitation." In *Imitation in Animals and Artifacts*, 1–40. Cambridge, MA: MIT Press.
- Dempster, Arthur P., Nan M. Laird, and Donald B. Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society: Series B (Methodological)* 39 (1): 1–22.
- Folgieri, Raffaella, and Roberto Zampolini. 2015. "BCI Promises in Emotional Involvement in Music and Games." *Computers in Entertainment* 12 (1): 1–10.
- Fujimoto, Scott, Herke van Hoof, and David Meger. 2018. "Addressing Function Approximation Error in Actor-Critic Methods." ArXiv preprint: 1802.09477.
- Geng, Tao, John Q. Gan, and Huosheng Hu. 2010. "A Self-Paced Online BCI for Mobile Robot Control." *International Journal of Advanced Mechatronic Systems* 2 (1–2): 28–35.
- Guler, Inan, and Elif Derya Ubeyli. 2007. "Multiclass Support Vector Machines for EEG-Signals Classification." *IEEE Transactions on Information Technology in Biomedicine* 11 (2): 117–126.
- Hermens, Hermie J., Bart Freriks, Catherine Disselhorst-Klug, and Günter Rau. 2000. "Development of Recommendations for SEMG Sensors and Sensor Placement Procedures." *Journal of Electromyography and Kinesiology* 10 (5): 361–374.
- Hu, Jianbing, Zining Yang, Zhiyang Wang, Xinyu Wu, and Yongsheng Ou. 2015. "Neural Learning of Stable Dynamical Systems Based on Extreme Learning Machine." In *2015 IEEE International Conference on Information and Automation*, 306–311. New York: IEEE.
- Jang, Eric, Coline Devin, Vincent Vanhoucke, and Sergey Levine. 2018. "Grasp2vec: Learning Object Representations from Self-Supervised Grasping." ArXiv preprint: 1811.06964.
- Jiang, Yiming, Yang Chenguang, Na Jing, Li Guang, Li Yanan, and Zhong Junpei. 2017. "A Brief Review of Neural Networks Based Learning and Control and Their Applications for Robots." *Complexity* 2017:1895897.
- Kalashnikov, Dmitry, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, et al. 2018. "Qt-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation." ArXiv preprint: 1806.10293.
- Khansari-Zadeh, S. Mohammad, and Aude Billard. 2011. "Learning Stable Nonlinear Dynamical Systems with Gaussian Mixture Models." *IEEE Transactions on Robotics* 27 (5): 943–957.
- Khezri, Mahdi, and Mehran Jahed. 2007. "Real-Time Intelligent Pattern Recognition Algorithm for Surface EMG Signals." *Biomedical Engineering Online* 6 (1): 1–12.

- Kosmyna, Nataliya, Franck Tarpin-Bernard, Nicolas Bonnefond, and Bertrand Rivet. 2016. "Feasibility of BCI Control in a Realistic Smart Home Environment." *Frontiers in Human Neuroscience* 10:416.
- Kumra, Sulabh, and Christopher Kanan. 2017. "Robotic Grasp Detection Using Deep Convolutional Neural Networks." In *Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 769–776. New York: IEEE.
- Lee, Michelle A., Yuke Zhu, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette Bohg. 2019. "Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks." In *2019 International Conference on Robotics and Automation*, 8943–8950. New York: IEEE.
- Levine, Sergey, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. 2018. "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection." *International Journal of Robotics Research* 37 (4–5): 421–436.
- Liang, Hengyue, Xibai Lou, and Changhyun Choi. 2019. "Knowledge Induced Deep Q-Network for a Slide-to-Wall Object Grasping." ArXiv preprint: 1910.03781.
- Lillicrap, Timothy P., Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. "Continuous Control with Deep Reinforcement Learning." ArXiv preprint: 1509.02971.
- MacQueen, James. 1967. "Some Methods for Classification and Analysis of Multivariate Observations." *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 1 (14): 281–297.
- Matas, Jan, Stephen James, and Andrew J. Davison. 2018. "Sim-to-Real Reinforcement Learning for Deformable Object Manipulation." ArXiv preprint: 1806.07851.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, et al. 2015. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518 (7540): 529–533.
- Morrison, Douglas, Peter Corke, and Jürgen Leitner. 2018. "Closing the Loop for Robotic Grasping: A Real-Time, Generative Grasp Synthesis Approach." ArXiv preprint: 1804.05172.
- Müller-Putz, Gernot R., Reinhold Scherer, Christian Brauneis, and Gert Pfurtscheller. 2005. "Steady-State Visual Evoked Potential (SSVEP)-Based Communication: Impact of Harmonic Frequency Components." *Journal of Neural Engineering* 2 (4): 123.
- Mülling, Katharina, Jens Kober, Oliver Kroemer, and Jan Peters. 2013. "Learning to Select and Generalize Striking Movements in Robot Table Tennis." *International Journal of Robotics Research* 32 (3): 263–279.
- Murray, Richard M., Zexiang Li, S. Shankar Sastry, and S. Shankara Sastry. 1994. *A Mathematical Introduction to Robotic Manipulation*. Boca Raton: CRC Press.
- Nair, Ashvin V., Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. 2018. "Visual Reinforcement Learning with Imagined Goals." *Advances in Neural Information Processing Systems* 31 (2018): 9191–9200.
- Osu, Rieko, David W. Franklin, Hiroko Kato, Hiroaki Gomi, Kazuhisa Domen, Toshinori Yoshioka, and Mitsuo Kawato. 2002. "Short-and Long-Term Changes in Joint Co-contraction Associated with Motor Learning as Revealed from Surface EMG." *Journal of Neurophysiology* 88 (2): 991–1004.
- Peng, Guangzhu, Chenguang Yang, Yiming Jiang, Long Cheng, and Peidong Liang. 2016. "Teleoperation Control of Baxter Robot Based on Human Motion Capture." In *2016 IEEE International Conference on Information and Automation*, 1026–1031. New York: IEEE.
- Quillen, Deirdre, Eric Jang, Ofir Nachum, Chelsea Finn, Julian Ibarz, and Sergey Levine. 2018. "Deep Reinforcement Learning for Vision-Based Robotic Grasping: A Simulated Comparative Evaluation of Off-Policy Methods." In *2018 IEEE International Conference on Robotics and Automation*, 6284–6291. New York: IEEE.
- Ray, G. C., and S. K. Guha. 1983. "Relationship between the Surface EMG and Muscular Force." *Medical and Biological Engineering and Computing* 21 (5): 579–586.
- Rebsamen, Brice, Etienne Burdet, Cuntai Guan, Haihong Zhang, Chee Leong Teo, Qiang Zeng, Christian Laugier, and Marcelo H. Ang. 2007. "Controlling a Wheelchair Indoors Using Thought." *IEEE Intelligent Systems* 22 (2): 18–24.
- Redmon, Joseph, and Anelia Angelova. 2015. "Real-Time Grasp Detection Using Convolutional Neural Networks." In *2015 IEEE International Conference on Robotics and Automation*, 1316–1322. New York: IEEE.
- Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. "You Only Look Once: Unified, Real-Time Object Detection." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 779–788. New York: IEEE.
- Rozo, Leonel, Pablo Jiménez, and Carme Torras. 2013. "A Robot Learning from Demonstration Framework to Perform Force-Based Manipulation Tasks." *Intelligent Service Robotics* 6 (1): 33–51.
- Schaal, Stefan. 1999. "Is Imitation Learning the Route to Humanoid Robots?" *Trends in Cognitive Sciences* 3 (6): 233–242.

- Schaal, Stefan, Peyman Mohajerin, and Auke Ijspeert. 2007. "Dynamics Systems vs. Optimal Control—a Unifying View." *Progress in Brain Research* 165:425–445.
- Siciliano, Bruno, and Oussama Khatib, eds. 2016. *Springer Handbook of Robotics*. Berlin: Springer.
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, et al. 2016. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529 (7587): 484–489.
- Trottier, Ludovic, Philippe Giguere, and Brahim Chaib-Draa. 2017. "Convolutional Residual Network for Grasp Localization." In *2017 14th Conference on Computer and Robot Vision*, 168–175. New York: IEEE.
- Tsarouchi, Panagiota, Sotiris Makris, and George Chryssolouris. 2016. "Human-Robot Interaction Review and Challenges on Task Planning and Programming." *International Journal of Computer Integrated Manufacturing* 29 (8): 916–931.
- Tsui, Chun Sing Louis, John Q. Gan, and Huosheng Hu. 2011. "A Self-Paced Motor Imagery Based Brain-Computer Interface for Robotic Wheelchair Control." *Clinical EEG and Neuroscience* 42 (4): 225–229.
- Wolpaw, Jonathan R., Niels Birbaumer, William J. Heetderks, Dennis J. McFarland, P. Hunter Peckham, Gerwin Schalk, Emanuel Donchin, Louis A. Quatrano, Charles J. Robinson, and Theresa M. Vaughan. 2000. "Brain-Computer Interface Technology: A Review of the First International Meeting." *IEEE Transactions on Rehabilitation Engineering* 8 (2): 164–173.
- Wu, Huaiwei, Chenguang Yang, Ning Wang, Wei He, and Chun-Yi Su. 2017. "Manipulation of a Robot Arm in 3D Space by Using EEG Signals." In *2017 2nd International Conference on Advanced Robotics and Mechatronics*, 608–613. New York: IEEE.
- Wu, Yue, Hui Wang, Biaobiao Zhang, and K-L. Du. 2012. "Using Radial Basis Function Networks for Function Approximation and Classification." *ISRN Applied Mathematics* 2012:324194.
- Xu, Yang, Chenguang Yang, Junpei Zhong, Hongbin Ma, Lijun Zhao, and Min Wang. 2017. "Robot Teaching by Teleoperation Based on Visual Interaction and Neural Network Learning." In *2017 9th International Conference on Modelling, Identification and Control*, 1068–1073. New York: IEEE.
- Yang, Chenguang, Chuize Chen, Wei He, Rongxin Cui, and Zhijun Li. 2019. "Robot Learning System Based on Adaptive Neural Control and Dynamic Movement Primitives." *IEEE Transactions on Neural Networks and Learning Systems* 30 (3): 777–787.
- Yang, Chenguang, Chuize Chen, Ning Wang, Zhaojie Ju, Jian Fu, and Min Wang. 2019. "Biologically Inspired Motion Modeling and Neural Control for Robot Learning from Demonstrations." *IEEE Transactions on Cognitive and Developmental Systems* 11 (2): 281–291.
- Yang, Chenguang, Yiming Jiang, Jing Na, Zhijun Li, Long Cheng, and Chun-Yi Su. 2018. "Finite-Time Convergence Adaptive Fuzzy Control for Dual-Arm Robot with Unknown Kinematics and Dynamics." *IEEE Transactions on Fuzzy Systems* 27 (3): 574–588.
- Yang, Chenguang, Jing Luo, Chao Liu, Miao Li, and Shi-Lu Dai. 2018. "Haptics Electromyography Perception and Learning Enhanced Intelligence for Teleoperated Robot." *IEEE Transactions on Automation Science and Engineering* 16 (4): 1512–1521.
- Yang, Chenguang, Jing Luo, Yongping Pan, Zhi Liu, and Chun-Yi Su. 2018. "Personalized Variable Gain Control with Tremor Attenuation for Robot Teleoperation." *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 48 (10): 1759–1770.
- Yang, Chenguang, Hongbin Ma, and Mengyin Fu. 2016. *Advanced Technologies in Modern Robotic Applications*. Singapore: Springer.
- Yang, Chenguang, Huaiwei Wu, Zhijun Li, Wei He, Ning Wang, and Chun-Yi Su. 2018. "Mind Control of a Robotic Arm with Visual Fusion Technology." *IEEE Transactions on Industrial Informatics* 14 (9): 3822–3830.
- Yang, Chenguang, Chao Zeng, Yang Cong, Ning Wang, and Min Wang. 2019. "A Learning Framework of Adaptive Manipulative Skills from Human to Robot." *IEEE Transactions on Industrial Informatics* 15 (2): 1153–1161.
- Yang, Chenguang, Chao Zeng, Cheng Fang, Wei He, and Zhijun Li. 2018. "A DMPS-Based Framework for Robot Learning and Generalization of Humanlike Variable Impedance Skills." *IEEE/ASME Transactions on Mechatronics* 23 (3): 1193–1203.
- Yang, Chenguang, Chao Zeng, Peidong Liang, Zhijun Li, Ruifeng Li, and Chun-Yi Su. 2017. "Interface Design of a Physical Human-Robot Interaction System for Human Impedance Adaptive Skill Transfer." *IEEE Transactions on Automation Science and Engineering* 15 (1): 329–340.
- Young, Brittany M., Zack Nigogosyan, Veena A. Nair, Léo M. Walton, Jie Song, Mitchell E. Tyler, Dorothy F. Edwards, et al. 2014. "Case Report: Post-Stroke Interventional BCI Rehabilitation in an Individual with Preexisting Sensorineural Disability." *Frontiers in Neuroengineering* 7:18.

- Yuan, Weihao, Kaiyu Hang, Danica Kragic, Michael Y. Wang, and Johannes A. Stork. 2019. "End-to-End Non-prehensile Rearrangement with Deep Reinforcement Learning and Simulation-to-Reality Transfer." *Robotics and Autonomous Systems* 119:119–134.
- Zeng, Andy, Shuran Song, Stefan Welker, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. 2018. "Learning Synergies between Pushing and Grasping with Self-Supervised Deep Reinforcement Learning." In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4238–4245. New York: IEEE.
- Zhang, Hanbo, Xuguang Lan, Xinwen Zhou, and Nanning Zheng. 2018. "ROI-Based Robotic Grasp Detection in Object Overlapping Scenes Using Convolutional Neural Network." ArXiv preprint: 1808.10313.
- Zhang, Jiahao, Chenguang Yang, Miao Li, and Ying Feng. 2018. "Grasping Novel Objects with Real-Time Obstacle Avoidance." In *International Conference on Social Robotics*, 160–169. Cham, Switzerland: Springer.
- Zhao, Jing, Wei Li, Xiaoqian Mao, and Mengfan Li. 2015. "SSVEP-Based Experimental Procedure for Brain-Robot Interaction with Humanoid Robots." *Jove (Journal of Visualized Experiments)* 105:E53558.

17

Cognitive Control for Decision and Human-Robot Collaboration

Erwin Jose Lopez Pulgarin, Ute Leonards, and Guido Herrmann

This chapter focuses on the concept of cognitive control in robotics and how it is linked to decision, control, and human-robot interaction (HRI). Achieving a control paradigm that enables robust, flexible goal-driven performance in a myriad of scenarios involving unstructured changing environments and interaction between robots and other agents such as humans has been pursued during the last decade (e.g., Avery, Kelley, and Davani 2006; Baud-Bovy et al. 2014; Herrmann and Leonards 2018). In order to achieve this, inspiration has been taken from nature, with a focus on the way humans and other animals undertake their decision and control processes (see chapter 1). Indeed, by creating controllers inspired by human flexibility and adaptability, some or all of the qualities found in human cognitive processes can be pursued (i.e., adaptability, robustness, goal-driven behavior with sensor and subtask prioritization) in artificial programmable systems.

First, this chapter includes an introduction to the concept of control in the context of industrial processes and expands it to robotics in general; challenges behind robot control will be raised, highlighting the need for novel decision and control architectures for modern robotics such as those involved in closely interacting with humans, dealing with unstructured environments, and learning to better perform a task—hence cognitive control.

Second, the word “cognitive” in the context of control will be defined after an overview about how “cognition” has been used in the literature; the definition of what a cognitive controller is will include aspects about both its architecture and inputs, highlighting how it relates to the term originally used in human behavioral studies and cognitive neuroscience.

Finally, a modeling approach for cognitive control, which integrates the principles of multiagent interaction into a decision-making (i.e., discrete and probabilistic) and control action (i.e., continuous and dynamic) framework will be proposed. This will be followed by a discussion around the framework’s elements and their wider impact in different areas of application, such as autonomous driving, teleoperation, and human-humanoid interaction.

17.1 Control in Robotics

When considering any system that interacts with the environment and manipulates it by any means or in any way, the concept of control needs to be considered. Starting from industrial control or process control (Ogata 2010), the main objective of “control” science

is to be able to manipulate one or several variables of a system and make them behave as one desires. Control problems can be described generically as either a trajectory-following problem (i.e., make a system's variable follow a set of values) or a regulation problem (i.e., keep a system's variable at a fixed value). Most modern control problems deal with closed-loop control architectures, using sensor or estimation inputs from the system to feed back to the controller; this feedback allows a comparison of expected system outputs with real outputs, which is a prerequisite to modify control outputs based on the state of the system (i.e., outputs). When considering the controller in a system description (Ogata 2010; Maciejowski 2002), it can be described based on its inputs (i.e., single input [SI] or multiple input [MI]) and its outputs (i.e., single output [SO] or multiple output [MO]), with its subsequent combinations (e.g., single input, single output [SISO] or single input, multiple output [SIMO] and so on). This relates to the system's complexity and the control goals—that is, the amount of inputs being how many sensor inputs or control goals the system requires and the amount of outputs as control signals or controlled variables. When considering the controller's inner workings, an explicit understanding of the system to be controlled is used and most of the time is needed in the form of a mathematical description of its dynamics. This understanding and the requirements for control determine how the controller inputs relate to the desired outputs. Based on the level of detail these models require, they could be described using any sort of mathematical description, such as linear operators, non-linear equations, and probability distributions, usually in a dynamic framework. Performance criteria are imposed on the controller in order to have a complete description of how each variable is controlled (e.g., time to reach the desired value, percentage of error when reaching the desired value, maximum error if the controller overshoots). Finally, controllers can be designed to deal with uncertainty from the system model and to be adaptable to changes in the environment or changes in the model itself. Figure 17.1 shows a general description of a control architecture, considering its required input (i.e., the system's demanded output), the controller that looks to achieve this input, the system, plant, or environment to be controlled, and the sensory input that comes from the system itself.

Bringing these concepts from industrial machinery to the realm of robotics was a straightforward task in the early stages of robotics, as industrial robots had similar physical shapes and objectives compared to industrial machinery (i.e., industrial manipulators were dealing with repetitive tasks with high precision at high speeds). Indeed, most industrial controller designs focused on dealing with low-level control for each link or motor, while

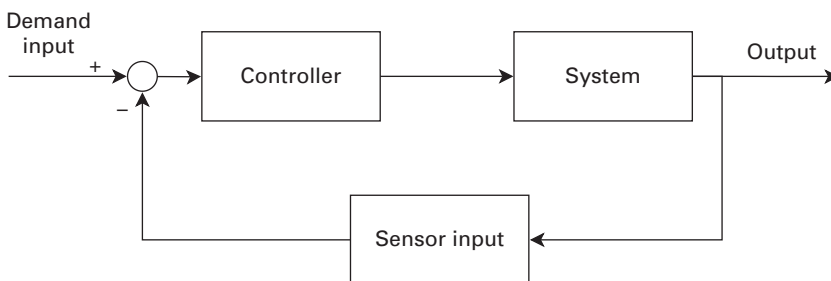


Figure 17.1
General control architecture.

high-level path and trajectory planning was dealt with through solutions based on the robot's geometric properties (e.g., a kinematic description using a Jacobian for end-effector positioning or forward kinematics and motion planning or inverse kinematics). These approaches were highly successful for a wide variety of industrial applications (see LaValle 2006; Scassellati 2002; Visioli and Legnani 2002).

However, as the field of robotics expands, the desire to move robots from industrial setups to more general environments brings challenges beyond what previous approaches can solve. First of all, the goals of a robot outside an industrial setup are potentially more generic and difficult to define completely in advance. Robots thus need to be able to change/adapt over time. For example, taking care of an elderly person could start with only checking their temperature and helping with mobility inside a room but might then evolve into reaching for objects, general social companionship, administering medicine, and more. In addition, using robots outside an industrial setup involves dealing with unstructured, complex, and changing environments that could be difficult to assess or predict at all times. Finally, some applications, such as robots for retailing, teaching, and medical care, would require interaction and/or cooperation with other autonomous agents—be it other robots or human beings (i.e., human-robot interaction). These are all challenges that go well beyond what traditional frameworks focusing on motion control would be able to deal with. Going back to the closed-loop controller description, any such system requires many multiple-input, multiple-output (MIMO) controllers with potentially nonlinear models, configured for both trajectory following and regulation just to focus on general movement alone—for example, to move the robot body to a known location, traverse unknown terrain, or mediate closeness to interacting robots or humans while maintaining safety. Additional components such as high-level decision-making and multimodal communication, supported by specialized hardware such as sensing, actuating, and communication devices, would be necessary to complement the proposed controller (Whitsell and Artemiadis 2017).

The goal here is to find an architecture or methodological approach that can help solve such problems in a complete and integrated manner. To achieve this, inspiration has been drawn from nature and, particularly, from human cognitive processes to better replicate and improve robots in “humanlike capabilities” such as dealing with unstructured and uncertain environments or prioritizing between subtasks and sensory input while maintaining a goal-driven task execution that is adaptable and changes over time. Indeed, human beings are the best-known system to date for adapting to new environments, performing robustly, and prioritizing while reaching a goal. In addition, it has been suggested that a robot that tries to copy or mimic human capabilities by relying on similar mechanisms as the person it is interacting with might be the easiest to understand intuitively (e.g., non-verbally) when interaction between artificial agents and humans is needed (Eder, Harper, and Leonards 2014).

17.2 Cognition in Control and Robotics

The use of the word “cognition” for control has been suggested because it takes inspiration from human cognitive processes. Cognition in humans covers mental processes and their role in thinking, feeling, and behaving, as defined by Kellogg (2015). Cognition includes

perception—that is, the processing and understanding of the outside world by sensory inputs (Fischer and Demiris 2019); *memory*, or how information is stored, manipulated, and used (Baddeley 2012); *decision-making*, or how to decide on the best action to reach a certain goal (Haefner, Berkes, and Fiser 2016); and *acquisition of knowledge and expertise*, including abstracting high-level understanding and learning from its interaction with the world (Moulin-Frier et al. 2018), among other factors, such as creativity and reasoning, as aspects of human capabilities.

17.2.1 Cognitive Architectures

In robotics, the concept of cognitive architecture comes from research in the field of artificial intelligence to describe a list of components, organizational structures, information flows, representations, and computational procedures that enable some intelligent behavior (Kotseruba and Tsotsos 2020; see also chapter 10); these mechanisms mimic ways the brain is thought to deal with and manipulate information. Such architectures tend to work as blueprints, with no consideration or explanation of how to be implemented in any specific agent. This means that they can be software based only or embodied in the form of a robot body (Kawamura et al. 2008; Wei and Hindriks 2013; see chapter 11). They focus on describing different “cognitive” modules that enable the mimicking of certain intelligent capabilities such as short- and long-term memory modules for better decision-making (Ratanaswad, Gordon, and Dodd 2005). Such modular descriptions tend to focus on the modules’ interconnections, their interaction with the outside world (i.e., environment) in the form of sensor inputs (i.e., stimuli), and their possible control outputs (i.e., action).

A wide range of cognitive architectures have been proposed over the past forty years, each author tackling the problem of representing humanlike intelligence or capabilities in their own way (see Kotseruba and Tsotsos [2020] for a recent review and chapter 10). A possible general classification for these architectures lies in the way information is processed and represented, either by using a handcrafted symbolic representation (i.e., symbolic or cognitivist systems), a sensor and data-based representation (i.e., emergent or connectionist systems), or a mix of both (i.e., hybrid systems; Kotseruba and Tsotsos 2020). Symbolic systems tend to have a long design process because they require a large initial knowledge base including rules, conditions, label descriptions, or possible scenario descriptions. They achieve great predictability and reproducibility, although at the expense of flexibility and robustness to changing environments. In contrast, emergent systems are highly adaptable, suited for learning from the environment and easier to design, but they require potentially long training processes, losing transparency in their results and traceability due to these learning processes. It thus becomes difficult to know what to learn, what exactly is being learned, and when to stop learning in order to achieve optimal performance.

The above classification serves as a parallel to one often employed in control science to describe the mathematics used to design and create the controller itself (Lopez Pulgarin et al. 2018): model-based controllers are designed using a mathematical representation of the system (i.e., plant model) that describes the dynamics surrounding the system. Such controllers are in stark contrast to a data-driven controller that uses available environment measurements to construct a relation between how a system is manipulated (i.e., actions) and the system itself (i.e., states) based on rewarding or punishing certain behaviors and

limited to no knowledge of the system itself (e.g., Al-Tamimi, Lewis, and Abu-Khalaf 2007; Na et al. 2012; Lewis, Vrabie, and Vamvoudakis 2012).

17.2.2 Cognitive Controllers

Cognitive controllers then are those that allow the creation of a controller by either implementing or taking inspiration from cognitive architectures (Haykin et al. 2012; Fatemi and Haykin 2014; Kawamura and Gordon 2006). Note that some authors define cognitive control as an addition to other low-level adaptive controllers (Haykin et al. 2012) or as a supplementary way to deal with high sensor input in parallel in a data-driven fashion while ignoring noncritical information (Kawamura and Gordon 2006). Yet, even for such alternative uses of the word “cognitive,” authors generally agree on the idea of drawing inspiration from mental models or brain-inspired cognitive architectures. As many cognitive architectures exist, however, there is no single standard of how the components should look (i.e., submodules, types of inputs or outputs, functionality implemented) and thus how these intelligent/mental capabilities are achieved. Figure 17.2 shows an adaptive controller, an extension of the architecture shown in figure 17.1, that allows the model to learn from the environment and inform the controller of some previously unknown parameters in the system to allow it to adapt (Khan et al. 2012; Na et al. 2015). In cognitive architectures, these capabilities are embedded in a cognitive action module, where information derived from perception inform the system how to learn and adapt to the changing and unknown environment.

The main difference between a modern or smart controller (Kawamura et al. 2008) and a cognitive controller is their flexibility in goal description. Although both include interaction with the environment via sensory input and actuation output, having some kind of memory of the environment and the interaction of the controller with it, the cognitive controller is not restricted to one particular task; it has the capability to translate information to other tasks and thus goes beyond initial requirements. In other words, cognitive controllers have the ability to go beyond an initial task definition in order to

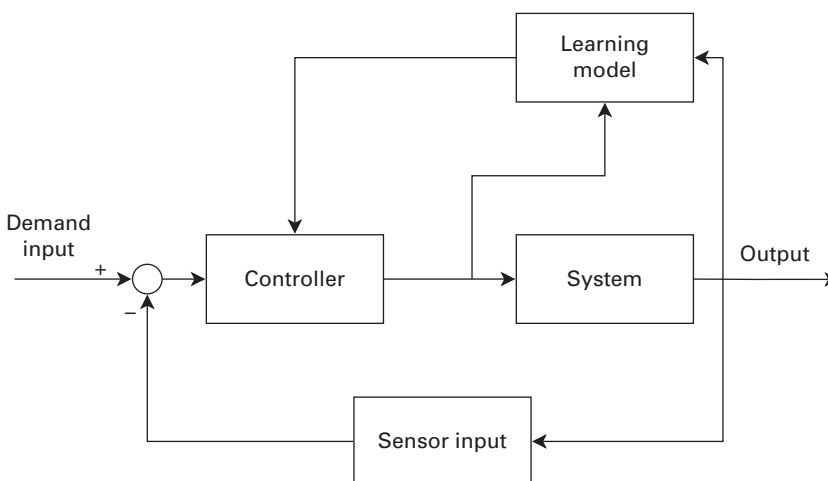


Figure 17.2
Adaptive controller architecture.

achieve an overarching goal through generalization and flexibility (Kotseruba and Tsotsos 2020).

Considering the goal of allowing high-level decision-making and control by a cognitive controller (Kotseruba and Tsotsos 2020), a more detailed cognitive architecture can be formulated by reviewing the specifics of human cognitive processes (Kellogg 2015; e.g., perception, memory, learning). Figure 17.3 introduces the general information loop used in many cognition-inspired applications (Kawamura and Gordon 2006; Ratanaswasd, Gordon, and Dodd 2005), expanding the previously introduced perception and action modules. Sensing and actuation are separated, suggesting they deal only with how sensory information is transformed into useful knowledge and information (i.e., perception) and how the selected decision or sets of actions are performed (i.e., actuation and low-level control), respectively (Haefner, Berkes, and Fiser 2016). A module is added that deals with both the regulation and control of how perception outcomes are used (Gold and Heekeren 2013) and how they can relate to a specific goal such as executive functions or more general goal-related information. An additional module (Ratanaswasd, Gordon, and Dodd 2005) is added that considers how all remaining modules can generate relevant information that could be stored and used to improve their functioning over time and how this process is performed (i.e., learning and memory); the inner workings of this module tend to take inspiration from working-memory models in humans (e.g., Baddeley 2000, 2012).

The information loop of decision-making and control in figure 17.3 implies that for a certain scenario the best possible decision is selected from any set of possibilities by cycling through them and performing any necessary motor control (e.g., limb movement, gaze control, speech). This loop resembles the problem faced in nonlinear control when

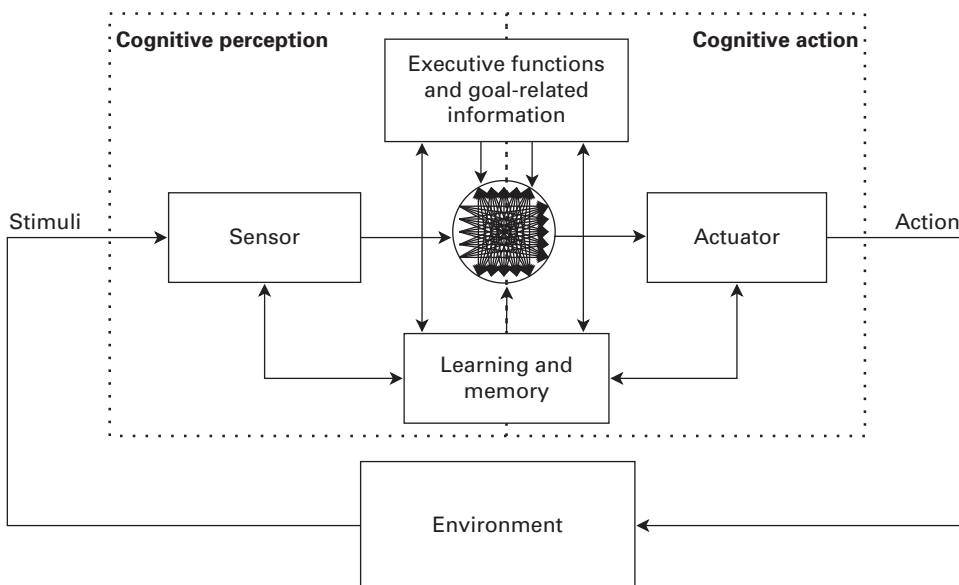


Figure 17.3

Cognitive control architecture with general functional blocks. *Source:* Inspired by Kazahiko Kawamura and Gordon 2006; Ratanaswasd, Gordon, and Dodd 2005.

dealing with uncertain or highly dynamic environments for which a certain controller has been specifically designed or tuned for optimal performance within a specific range of the dynamics, called *gain scheduling* (Yang et al. 2010). The challenges faced in gain scheduling could be seen as a reduced set of those arising in cognitive control: in the former, the cases for which a set of controllers is designed and the controllers themselves are known in advance, and the challenge is to tune the controllers and change from one to the other to maintain performance and stability; in the latter, an additional challenge is to select from an only vaguely defined set of uncertain possibilities and to perform control over them with little to no prior knowledge.

17.2.3 Control in Cognitive Robotics and HRI

Cognitive robotics (Levesque and Lakemeyer 2008; see also chapter 1) arises with the use of cognitive architectures or concepts inspired by these architectures in order to tackle challenges faced in robotics at both task (e.g., object manipulation, exploration) and application levels (e.g., autonomous operation, teleoperation, HRI), respectively. Tasks that have been performed in cognitive robotics range from command following for object manipulation (e.g., Ratanaswasd, Gordon, and Dodd 2005; Dodd and Gutierrez 2005; Kawamura and Gordon 2006; Kawamura et al. 2008) to autonomous navigation (e.g., Avery, Kelley, and Davani 2006; Wei and Hindriks 2013) to reaching a goal by changing tasks (e.g., Khamassi et al. 2011).

Building on such achieved robotic capabilities (e.g., object reaching and navigation), applications that go beyond following an explicit human command have been proposed that tend to involve humans in some aspect or another (e.g., medical aid; Neerinx et al. 2019); hence, Human-robot interaction (HRI) is involved. HRI is the term used to include all the tools and studies around the actuation and interaction of robots with human beings in any possible way (see also chapter 19). Cognitive robotics has proposed a range of methodologies to better interact with humans, such as knowledge and skill transfer from human to robot (e.g., Tan and Liang 2011), knowledge acquisition and learning through interaction (e.g., Moulin-Frier et al. 2018; Nakamura, Nagai, and Taniguchi 2018), and perspective taking (Fischer and Demiris 2019), to name but a few. However, robots with full autonomy have not yet been achieved.

Building from the definition of HRI, a special category focused on scenarios in which robot and human work together to reach a common goal is called human-robot collaboration (HRC). Two key methodological aspects of HRC highlighted by Bauer, Wollherr, and Buss (2008) in their review of the most challenging aspects of HRC are *intention* and *action*; the former considers an initial agreement of the common/joint goal either by explicit (e.g., speech and haptic commands) or implicit (e.g., hand gestures, eye gaze, estimation from physiological signals) means, and the latter considers planning and replanning capabilities to deal with unstructured dynamic environments and a potential joint action (e.g., carrying and sharing a moving load).

HRI brings challenges beyond those previously stated. Even if cognitive processes could be mimicked to better deal with an unstructured and uncertain environment following a certain goal, the challenge of interacting with an autonomous agent who deals with a similar cognitive architecture that requires dynamic change and adaptation is a daunting task. As human beings can perform many different tasks and actions with no guarantee

that they will do what the interacting robot expects, robots need to be equipped with the ability to both predict human actions effectively and to clearly communicate their intentions to the interacting human (e.g., Scassellati 2002; Grigore et al. 2013; Eder, Harper, and Leonards 2014; Herrmann and Leonards 2018).

17.3 A Multiagent-Inspired Approach to Control in Cognitive Robots

After having introduced cognitive robotics and its challenges, particularly for advanced HRI applications, we now move on to a decision and control action scheme (DCAS) that provides a clear application framework in which we try to tackle some of the issues raised above. This framework is focused on applications in which spatially close interaction or cooperation between human and robot is either a necessity or would at least improve overall task performance (e.g., semiautonomous vehicles or robotic care). The main challenge in these applications is to achieve safe, cooperative, human-centered, and human-predictive decision-making between a technological robotic device and a goal-oriented human through intelligent control and decision-making.

17.3.1 Paradigm Proposal for a Multiagent-Inspired Dynamic Decision and Action Framework for Human-Robot Interaction

Current state-of-the-art HRI sees the human as “in the loop” and thus as an unpredictable part of the robot’s cognitive control system (see, e.g., Eder, Harper, and Leonards 2014). The addition of the human inside a control loop means trying to model the human’s requirements, needs, or general behavior in order to minimize any negative effect on task performance or any risk of harming the human in close proximity to the robot while the robot navigates an environment (Dondrup et al. 2015). The uncertainty that arises from the “unpredictable” human can be dealt with safely and reliably as long as the environment in which such interactions happen is well controlled (Eder, Harper, and Leonards 2014). However, problems arise as soon as the environment itself becomes unpredictable. For most everyday environments, this is the case because they often include both other humans and animals (i.e., autonomous agents), making the environment unpredictable and demanding the system to interact or coordinate not only with one unpredictable partner but, potentially, with a whole range of external agents at the same time. Moreover, many physical environments themselves are too complex to be predicted in their entirety, thus leaving further risk of unpredictability. This means that we have an unpredictable part within the system itself as well as an unpredictable, continuously changing environment, a problem that is very hard to solve.

One way to solve this issue is by changing how one understands the directly collaborating partner and their role relative to the robot. If we understand the robot as an autonomous yet collaborative agent in its own right and take the human out of its direct loop by understanding them as an autonomous partner in the robot’s environment, then we have to solve only one issue—namely, the dynamic environmental uncertainty or unpredictability. As a partner, the human has built an internal model of the autonomous agent (e.g., robot or another human), as much as the autonomous agent has an internal model regarding the human colead/any other human in the environment. In cognitive psychology terms, such

an internal model of an interaction partner’s mind would be based on a concept known as *theory of mind* (Baron-Cohen et al. 1985). Theory of mind refers to the attribution of mental states (e.g., intentions, beliefs, and desires) to living beings; for an interaction scenario between two people, an understanding of the other agent’s intentions and decision-making process is essential for seamless interaction. Translated to HRI, there is thus only an “intensity” proximity difference or connectivity between the human and other autonomous agents in the environment, comparable to human-human interaction in close proximity or further away (i.e., personal space or extrapersonal space; Curioni, Knoblich, and Sebanz 2017). Hence, we suggest a scenario in which an autonomous system and a human each act as independent autonomous agents. As in human-human interaction, the two interacting partners can then have substantially different abilities as long as their internal representation of each other is sufficiently accurate.

This creates a redundant, safe, and interchangeable cooperative dynamic partnership between the “lead” and “colead” in which both robot and human can take on either role (Curioni, Knoblich, and Sebanz 2017). Communication and cooperation between the autonomous system and the human are a necessity not only for safety reasons but also for the accomplishment of common objectives as determined by the human. The joint action process between an artificial agent and a human being can only realize the optimal outcome of safe and efficient cooperation (i.e., shared control) if the autonomous system is able to synthesize, evaluate, and predict the human colead’s intentions and communicate its own possibly limited aims and capabilities to the interacting partner and the environment more generally (figure 17.4). This can be achieved as a cooperative decision and a subsequent dynamic

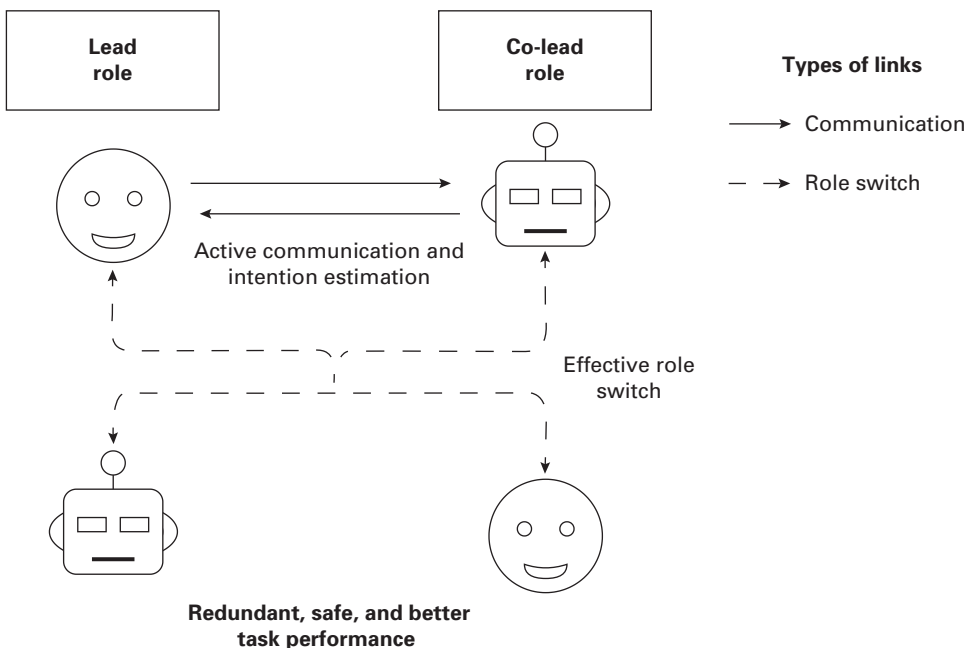


Figure 17.4

Autonomous robot: roles and their interchange between a human and an autonomous system.

action. In fact, such communication and cooperation are key, yet highly problematic, for HRI in general (Herrmann and Leonards 2018). The following suggestion of a decision-and-action framework provides a possible basis for a technical, dynamic HRI control paradigm to deal with interaction issues.

The above proposed paradigm shift in the way we think about fast dynamic interactions between people and artificial autonomous systems (i.e., robots) looks at the interaction and cooperation of two cooperative autonomous agents (figure 17.5) who operate as an interchangeable lead and colead (figure 17.4). Both agents are engaged in the task, and any inattention or objective track loss can be detected immediately. We propose a fluent change between who leads and who follows in joint actions in line with what is known for human-human interaction (Curioni, Knoblich, and Sebanz 2017). Indeed, coordination with others is implicit in many of our human behaviors. Such principles of cooperation can be nicely framed in a theoretical cooperative hybrid decision and dynamic control framework, the technical instantiation of the paradigm shift in dynamic HRI.

We propose that the solution to any human-artificial agent interaction lies in creating an intelligent cooperative decision and actuation framework in which decision-making-relevant information can be seamlessly merged with the human's goals and interests through theory of mind, to the extent necessary and possible. Similar to human-human joint action, the autonomous agent becomes a partner in its own right that is jointly involved in the decision-making process.

Within this cooperative framework, it is important for each agent to be aware that there are other, possibly less capable, autonomous agents in the environment. In this development context, the autonomous system-human relationship can be seen as the pupil (robot)-teacher (human) relationship in a learning stage, with a relationship of a close set of trusted partners as the end goal. The willing and supportive autonomous agent learns how to better interpret and interact (i.e., the autonomous agent learns from and adapts to the human agent). There is also the need for a "human-agent-detection" method to pick up on "error signals" induced during a task (e.g., inattentiveness within the teacher) so corrective actions can be made.

The successful interaction between human and autonomous agent would have to be fluid. This requires both cooperative decisions and cooperative dynamic actions to guarantee a safe and trusted cooperative process during the decisive changeover of leader and follower. For such a technical mechanism of cooperative interaction between two autonomous agents to work, the guiding principle that underlies this cooperation needs to be based on optimality, a principle well known in engineering (Turnbull et al. 2016) and robotics (Mombaur, Truong, and Laumond 2010; Khan et al. 2012) as well as an underlying concept to cognitive science (Berkes et al. 2011; Fiser et al. 2010), where it has been shown that under most circumstances humans decide and dynamically act in an optimal sense (e.g., Spiers, Khan, and Herrmann 2016; Haefner, Berkes, and Fiser 2016).

Putting the different concepts together, a hybrid optimal, yet adaptive, cooperative agent-based decision and control action scheme (i.e., DCAS) must provide the "intelligence" as an active negotiation scheme between autonomous agent and human. This scheme must resolve both the dynamic, the physical, and the behavioral event-driven interaction between human and autonomous system. To date, this is still an important unresolved step.




Figure 17.5
Human-autonomous agent interaction principles, appropriated from human-human interaction.

17.3.2 Principles and Characteristics of the Dynamic Decision and Action Framework

Based on predictions of the possible decisions a human agent could make (Lopez Pulgarin, Herrmann, and Leonards 2018), any DCAS should look at making decisions within a fraction of those prediction windows (e.g., one second) to then dynamically cue and actively influence the decisions of the human partner. Hence, human and autonomous agent would be able to cooperatively act within the time period of the predicted decisions and actions thereafter. The following axioms would lead to the DCAS:

- 1) The realization that we can treat the human as an “external source” or independent collaboration partner in relation to the autonomous agent instead of “in the loop.”
- 2) Learning from and adapting to the human and the signals people send in joint action situations, proxemics, and so on; learning to understand and predict adaptability within the human and their trust of the autonomous system as an autonomous, collaborative agent or partner.
- 3) Identification and subsequent learning from the “error” signals when situations go wrong. The principle of optimality of decisions and actions in human agents and control technology is an exploited commonality.
- 4) Identification and enabling of verbal and nonverbal communication channels in a human to indicate changes in “who is the leader, who is the follower” in joint action.
- 5) Subsequent “joint” cooperative, agent-based decision-making and dynamic action taking.

Overall, a coherent modeling methodology for decisions and actions would have to be developed that is deeply rooted in complementary research on human decision-making, cognition, communication, dynamic actions, dynamic decisions, and action theories in control and computer science.

This requires that agent models and their uncertainties involved in the joint decision-making process be predetermined. This includes both the human and the autonomous system. The more the autonomous system relies on principles that underlie successful human-human interaction, the easier it will be for the human to develop a theory of mind of the robot. Only an approach that allows the human to intuitively understand the “mind” of the robot and that takes into account that an agent’s own actions influence other agents’ actions and vice versa will make joint actions among intelligent autonomous systems and humans possible (King, Rowe, and Leonards 2011).

Autonomous artificial agent models take inspiration from the fact that human decision models (e.g., Bellet et al. 2009; Berkes et al. 2011) have strong similarity to discrete hybrid stochastic automata (DHSA; Bemporad and Di Cairano 2005). There is a decision-making level that is responsible for the decisions, resulting in subsequent dynamic actions at the automatic level. Hence, the decision-making level may imply a set of discrete yet uncertain decisions, each followed by an uncertain dynamic action. Decisions are carried out within a fraction of a second, while dynamic actions can extend over intervals of several seconds.

The probabilistic approach for the analysis of human decision-making based on Fiser’s sampling-based probabilistic representational framework (Haefner, Berkes, and Fiser 2016; Fiser et al. 2010) is a possible guidance for the development of such agent models. In

Fiser’s framework, both the human’s internal representation of visual, aural, and tactile events during acting as a colead and the decision-making process in lead situations must be assessed. For the sequential character of decisions and dynamic actions, it therefore becomes necessary to explore how decisions in the present moment depend on the series of decisions made in the recent past. This leads to an assessment process of cues given to the human and the decisions made. For modeling the human decision-making process, the optimality principle following a Bayesian method can be used, such as the “cognitive tomography” method of Houlby et al. (2013). Applied to behavioral tasks, this allows for a quantitative description of an internal representation of a human based on discrete test choices (figure 17.6). Alternatively, a machine-learning-based understanding of the decision-making model (Lopez Pulgarin, Herrmann, and Leonards 2017) could be deployed and the synergies explored in which decision probabilities determine decision costs. Though such methods resemble emergent methods in cognitive architectures, they aim at presenting their results in a clearer and more predictable manner than traditional data-driven methods.

For the lower automatic dynamic action level—that is, the dynamic action following the decision—learning-based, regressive models based on data-driven methods might be preferable to strongly physical model-based methods; they may provide a continuous integral or summative optimal cost function that the human follows. Optimal cost function models allow for a more flexible prediction of the human’s actions. This is, for example, used in inverse optimal reinforcement learning (Mombaur, Truong, and Laumond 2010). Both levels are joint via the DHSA (Bemporad and Di Cairano 2005) and exploit mechanisms




Figure 17.6 Probabilistic internal decision model of a driver attempting to pass a car in front of them. *Source:* Adapted from Berkes et al. 2011.

like model predictive control (Morari and Lee 1999; Di Cairano et al. 2014; Rosolia, Zhang, and Borrelli 2018).

As mentioned earlier, a joining principle in human decision, human action dynamics, and many artificially designed technological processes is *optimality*. Each decision and action can be quantitatively associated with a cost. For a robot, a part of the cost in a dynamic action can be characterized, for instance, by its distance, either to a target for target tracking or to the distance from a human for safety. For both humans and robots, their energy consumption could be included in the cost and would be expected to increase over time while remaining limited in order for it to be optimized. The optimization of energy consumption underlies many human functions, such as locomotion (e.g., Warren 2006).

In terms of decision-making, the synergetic power of cognitive-science-founded models (e.g., Fiser et al. 2010; Berkes et al. 2011; figure 17.6) and machine-learning models (e.g., Lopez Pulgarin, Herrmann, and Leonards 2017, 2018) has to be exploited. Humans develop an internal model for each perceptive decision that guarantees that the decision regarding an intended task is carried out with the highest probability of success (Fiser et al. 2010; Berkes et al. 2011) considering the uncertainty of the environment (figure 17.7). Hence, decision costs are inversely related to the probability of the decision made. Identifying not only the models and their uncertainty sources but the optimal criteria for joint action between agents is key (Fiser et al. 2010).

The cooperative decision-making process can use the set of aforementioned DHSAs within a cooperative agent-based process, using model predictive control principles, to speed up the decision process and to allow fast computation of dynamic control actions from the multiagent framework. A probabilistic decision framework would possibly enhance such a process (Turnbull et al. 2016). For this, a virtual autonomous agent (figure 17.7) can

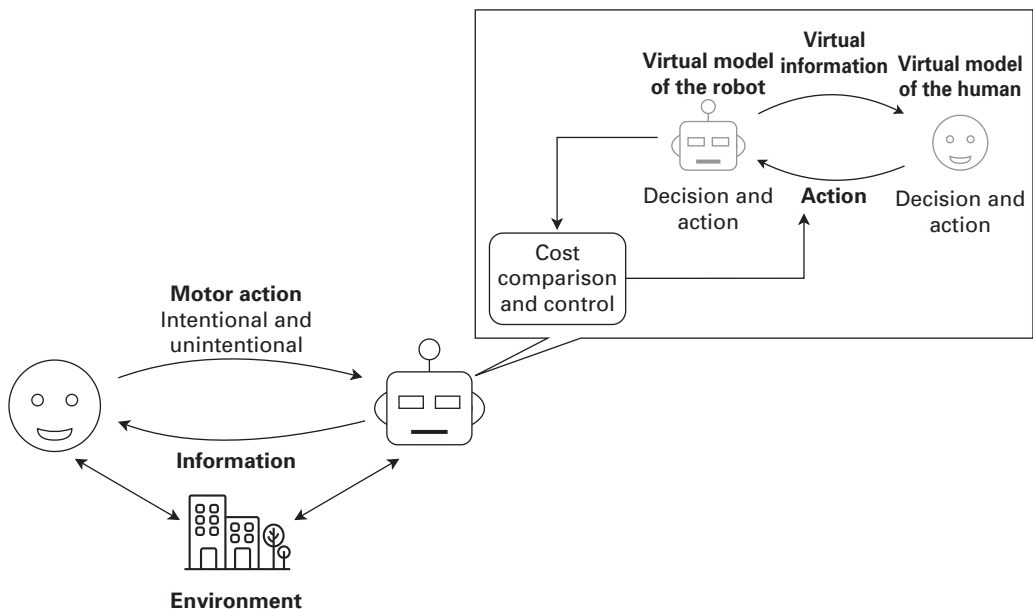


Figure 17.7
DCAS overview.

be developed by applying the principles behind DCAS. This virtual agent represents the nominal action computed from a joint optimal criterion for safety and a nominal understanding of the human, the internal model of the virtual leader, including individual differences between humans and in their intentions. The virtual model will act as an agent to be compared with the human characteristics and its short-term predictions using the unconstrained human model. Hence, the human and the virtual agent are assessed for their cost function, which evaluates whether the human cooperative partner is in line with the virtual human model. Subcomponents for safety are prioritized in decision-making, together with representations of human intention to decide to what extent the human or the virtual agent lead within the collaboration within a network of decentralized agents. Principles of game-theoretic approaches and agent synchronization can be used for a control policy in the vital time frame of dynamic actions following a decision, thus leading to an action of the cooperative decision-making process. To minimize conflicts, the autonomous system will take the human's desired actions as long as they do not compromise safety.

17.3.3 Impact on Autonomous Systems and HRI

Below we will analyze cognitive control for HRI with different types of robots and human-robot collaboration scenarios.

Humanoid robots

Humanoid robots (Oh, Kim, and Kim 2005) and interaction with them (i.e., human-humanoid interaction (HHI; Herrmann and Leonards 2018) could be enabled or improved by implementing a DCAS similar to that described above. Robots performing tasks that benefit from understanding the interacting human(s) while aiming at a final goal, such as to jointly move an object, to keep the human safe, or to maintain a human's vital signs inside a desired threshold, are the key benefits of the DCAS. Similar to existing cognitive architectures that aim at achieving humanlike capabilities, DCASs would allow many robotics applications to improve human life.

Understanding the interacting human and having the ability to share certain goals would be a big step toward safe, trustworthy HRI. For example, applications in medical assistive robotics could range from robots serving as partial nurses or assistants to medical professionals to shared physical cooperative work (e.g., object carrying; Parker and Croft 2012) or object manipulation (Sheng, Thobbi, and Gu 2015; Whitsell and Artemiadis 2017). By understanding the final goal that both the robot and medical professional share, meaning patient care, auxiliary actions could be performed by the robot across the whole care experience.

For cases in which the human is the recipient of the robot's actions and not the cooperative leader or companion, substantial benefits would be derived from understanding the human recipient's mindset in order to take the appropriate decisions at the best time possible.

Although the DCAS's main goal is not restricted to better understanding a robot's surrounding environment, it is one of its planned capabilities. Hence, the DCAS should improve the robot's autonomy during its sensing and decision-making processes by means of a collaborative learning strategy (e.g., supervised learning). By learning from the sensed environment while keeping a preset goal, long-term goals can be achieved autonomously and cooperatively as decision-making is improved across task iterations.

Teleoperated robots

Robotic teleoperation, understood to be the operation of a robot at a distance that allows one or many operators to interact with an environment (Li, Xia, and Su 2015), can benefit from the use of DCAS. As the scope of both operation and distance in teleoperation can be very wide (e.g., operation being by direct control or control by commands and distance understood as either a physical distance or difference in scale), many applications include a teleoperation setup (e.g., robotic surgeon, robotic manipulator for maintenance).

As in other HRI examples, DCAS would improve interaction to achieve a shared goal. Even if teleoperated robots are not considered autonomous or able to make decisions, the robot could possess intelligent mechanisms to help improve overall task performance—for example, to deal with potential delays in communication channels or complications introduced by control means or interfaces. By considering the robotic teleoperation device as a cooperative agent that understands and predicts the human operator's actions, the impact of delays could be minimized, as shared control would be made possible. This has been proposed before (e.g., Corredor, Sofrony, and Peer 2017), but here the idea is applied to a multitask and multidimensional space. Following a paradigm of a shared control, the level of autonomy in teleoperation devices could increase with improved understanding of the teleoperation task and increased safety.

In particular, higher autonomy of the system could speed up the operator's learning curve to use the device. Learning curve theory started empirically in the 1930s as cost reduction due to repetitive procedures in production plants was observed (see Anzanello and Fogliatto [2011] for the full reference); its goal is to exemplify and track how proficiency in performing a task or in the use of a device is improved via repetition (i.e., experience). Learning curves have been applied in teleoperation (e.g., Anvari 2007) to evaluate how much training is needed with using a device to achieve proficiency (Doumerc et al. 2010). Learning curves have been used in the field of medicine, particularly to evaluate both manual surgical procedures (e.g., Hopper, Jamison, and Lewis 2007; de Oliveira Filho 2002) and robotically assisted surgical procedures (e.g., Kaul, Shah, and Menon 2006; Chen et al. 2017) and to compare the two types of procedures with each other.

Building on the results around learning curves for robotic teleoperated devices, particularly in medicine (e.g., Yamaguchi et al. 2015; Samadi et al. 2007), a general learning curve can be proposed. Figure 17.8 shows the potential shape of the learning process behind a robotic device when plotting performance against experience. Three different phases can be identified: 1) an initial slow learning phase in which the operator gets used to the device until it reaches some minimal proficiency pg_1 after certain experience tp_1 , 2) a second practicing phase in which an acceptable proficiency pg_2 is achieved after continuous training tp_2 , and 3) a mastery phase in which optimal performance pg_3 is reached with continuous training and repetition.

DCAS could reduce training times tp_1 and tp_2 by making the teleoperation device both more intuitive and more responsive to the operator's needs. In addition, the gap between pg_1 and pg_2 could be reduced following the principle previously explained, ultimately leading to improvement in overall performance (i.e., push pg_3 higher).

The training of operators is an important task of teleoperation devices when autonomy levels of the teleoperation system are low. However, as the autonomy of a teleoperated robot increases, following autonomy levels similar to those declared by the Society of




Figure 17.8
Potential learning curve for the teleoperation of robotic devices.

Automotive Engineers (SAE; SAE International 2016), the DCAS could be an important enabler of improved teleoperation. Indeed, in many respects a teleoperated task is similar to a vehicular driving task in which increasing autonomy is introduced for improved performance, decreased human operator workload, and, ultimately, higher levels of safety.

Autonomous vehicles

Autonomous vehicles are a key target of many companies, as they could potentially bring significant economic and societal benefits (Fagnant and Kockelman 2015). Enormous structural efforts have been undertaken in terms of legislation and technology to enable autonomous driving. This includes the introduction of high-bandwidth G5 communication technology as an important enabler of autonomous driving through connectivity between cars or for high-precision maps. At the same time, the diverse and historically grown character of cities poses a challenge in its own right, with partially outdated infrastructure, differences in road regulations, and a highly dynamic environment due to other road users.

Albeit error prone, humans are fully capable of steering around a city's complexities. They can interpret complex situations, make decisions, resolve problems, and even reinterpret rules and road regulations within new contexts. Autonomous vehicles fail in such situations (see, e.g., fatal accidents with regard to Uber and Tesla; Banks, Plant, and Stanton 2018), meaning the human needs to remain included in the driving process. In addition, a significant number of countries, especially within Europe, demand a human-focused approach that requires the driver to be able to retake control at any moment—something that is not possible if a person has been occupied with a different task.

However, not only autonomous cars make mistakes. Drivers can be expected to make mistakes commensurate with the cognitive load they have to deal with or when they lack situational awareness through distraction or mind wandering (de Winter et al. 2014). While some advanced driver assistance systems (ADAS) and semiautonomous driving technologies try to account for human inattentiveness (e.g., Fagnant and Kockelman 2015), the majority work independently. Yet a more alert and experienced driving partner and copilot would be able to help steer a driver out of a temporary problem by direct communication or by providing supportive and intuitive cues for the driver.

Whether to allow human passengers to interact with autonomous cars remains an unresolved problem affecting cockpit design (Fagnant and Kockelman 2015), in addition to

the aforementioned uncertainty of understanding the human within the vehicle as part of the car's system (i.e., the human in the loop) and separating it conceptually from the external environment.

The DCAS suggested in this chapter interacts with the human in the car (i.e., the driver) in a cooperative way (figure 17.5), like a human pilot would with their human copilot. As pilot and copilot swap roles, so do the artificial agent (i.e., autonomous car) and the human driver, considering the requirements at hand and allowing the human agent to retake control of the driving process if desired.

17.4 Conclusion

A DCAS was introduced as a response to some of the challenges faced in modern robotics, such as goal-driven task performance and flexible and robust interaction with autonomous agents and the environment, as well as learning and knowledge acquisition. This decision and control framework was inspired by cognitive architectures and is expected to benefit many fields of application inside and beyond robotics. A list of a DCAS's major capabilities would be to

1. enable the robot's interaction with humans by understanding the human's goals and current state,
2. provide an agent-based description for both human and robot in order to enable joint action or cooperative work,
3. deal with partial or incomplete representations of the environment and the interacting agents using learning, and
4. exploit commonalities of recent research in human decisions and actions and existing predictive decision and action methodologies in control and decision theory.

However, many aspects of such a DCAS remain open questions, specifically of how to implement a cohesive mathematical framework around each of the scheme's components or capabilities. Going back to the previous list, some of its key challenges are as follows:

1. Human state and intention estimation and prediction
 - 1.1. What measurements can we use to help estimate or predict a human intention related to a certain task?
 - 1.2. How do we generate estimations or predictions of a human before, during, or after a task is being performed?
 - 1.3. How do we keep track of these estimations or predictions and update them as a task is being performed?
2. Task performance and coordination
 - 2.1. How do we make the robot perform a certain task or part of it?
 - 2.2. How do we let the robot know when to stop performing the task?
 - 2.3. How do we make the robot stop performing the task and release partial or complete control over a task?
 - 2.4. How do we let the robot know when to take back partial or full control of the task?
 - 2.5. How do we make the robot take back control of the task?

3. Decision-making and action with incomplete models

- 3.1. How do we integrate a learning process in a decision-making and control application?
- 3.2. How do we learn from performing a task and interacting with a human?
- 3.3. How do we convert sensed data and the learning process into knowledge useful for task completion and goal reaching?

Some technical insight has been given into how to answer these questions. A data-driven approach, taking advantage of both machine-learning (e.g., Lopez Pulgarin, Herrmann, and Leonards 2017, 2018; Khamassi et al. 2011) and probabilistic-sampling techniques (e.g., Nakamura, Nagai, and Taniguchi 2018; Haefner, Berkes, and Fiser 2016; Fiser et al. 2010), has been proposed as a feasible solution to improve understanding of the environment and to create knowledge, acknowledging challenges around modeling and validating and integrating the proposed methods into a more general cognitive control framework. Discrete hybrid automata (e.g., Bemporad and Di Cairano 2005) and model predictive control (e.g., Morari and Lee 1999) have been proposed as solutions for handling several action paths simultaneously (i.e., decision-making) and implementing controllers, with some others using reinforcement learning (i.e., data-driven methods) to deal with both situations (e.g., Lopez Pulgarin et al. 2018; Haykin et al. 2012; Khan et al. 2012; Khamassi et al. 2011). Hence, a suggested major joint guiding principle of these methods is optimality in discrete decisions and dynamic actions for dynamic autonomous agent-based cooperation. Some authors have managed to integrate data-driven methods with dynamical systems for control (e.g., Warren 2006), which again keeps the discussion going about how to better achieve a cognitive controller that takes advantage of symbolic (i.e., model-based) and emergent (i.e., data-driven) representations in cognitive architectures for control.

After introducing the concept of cognitive control and cognitive robotics, including its benefits and challenges, we hope to have sparked more interest in this promising research field while sharing some ideas and concepts developed over the past few years.

Acknowledgments

We would like to acknowledge the enormous contributions given by the following people in the form of discussions and idea sharing, which shaped the concepts described in this document. We would like to thank, in alphabetical order, Murad Abu-Khalaf, Eric Armengaud, Phil Barber, Gabriel Baud-Bovy, József Fiser, Tobias Kessler, Alois Knoll, Weiru Liu, Majid Mirmehdi, Henrik J. Putzer, Francesco Rea, Arthur Richards, Markus Rickert, Giulio Sandini, Alessandra Sciutti, and Robert Wragge-Morley.

Additional Reading and Resources

- An interesting book with applied examples of controllers for robotic arms movement: Spiers, Adam, Said Ghani Khan, and Guido Herrmann. 2016. *Biologically Inspired Control of Humanoid Robot Arms*. Cham, Switzerland: Springer.
- A comprehensive overview of some of the challenges in human-humanoid interaction inspiring work in cognitive robotics: Eder, Kerstin, Chris Harper, and Ute Leonards. 2014.

“Towards the Safety of Human-in-the-Loop Robotics: Challenges and Opportunities for Safety Assurance of Robotic Co-workers.” In *23rd IEEE International Symposium on Robot and Human Interactive Communication*, 660–665. New York: IEEE.

- A specific overview on optimal control and reinforcement learning, some of the techniques used in advanced control applications: Khan, Said G., Guido Herrmann, Frank L. Lewis, Tony Pipe, and Chris Melhuish. 2012. “Reinforcement Learning and Optimal Adaptive Control: An Overview and Implementation Examples.” *Annual Reviews in Control* 36 (1): 42–59. <https://doi.org/10.1016/j.arcontrol.2012.03.004>.
- ROS packages for symbolic planning and robot task planning: <https://moveit.ros.org/>, <http://wiki.ros.org/smach>, <http://wiki.ros.org/flexbe>.
- Software packages to get started with data-driven control (RL):
 - MATLAB (proprietary but with better documentation): <https://uk.mathworks.com/products/reinforcement-learning.html>.
 - PYTHON (free and more popular) for algorithms: <https://github.com/openai/baselines>; testing environments: <https://github.com/openai/gym>; use with robotic simulators http://wiki.ros.org/openai_ros.
- Software packages to get started with traditional control and model-based control (MPC):
 - Optimization solver: <https://osqp.org/>.
 - MATLAB (proprietary but with better documentation) control toolbox: <https://uk.mathworks.com/products/control.html>; MPC toolbox: <https://uk.mathworks.com/products/mpc.html>; modeling language wrapper: <https://yalp.github.io/>.
 - PYTHON (free) control library: <https://python-control.readthedocs.io/en/latest/>; free modeling language wrapper: <https://www.cvxpy.org/>.

References

- Al-Tamimi, Asma, Frank L. Lewis, and Murad Abu-Khalaf. 2007. “Model-Free Q-Learning Designs for Linear Discrete-Time Zero-Sum Games with Application to H-Infinity Control.” *Automatica* 43 (3): 473–481. <https://doi.org/10.1016/j.automatica.2006.09.019>.
- Anvari, M. 2007. “Remote Telepresence Surgery: The Canadian Experience.” *Surgical Endoscopy and Other Interventional Techniques*. Berlin: Springer. <https://doi.org/10.1007/s00464-006-9040-8>.
- Anzanello, Michel Jose, and Flavio Sanson Fogliatto. 2011. “Learning Curve Models and Applications: Literature Review and Research Directions.” *International Journal of Industrial Ergonomics* 41 (5): 573–583. <https://doi.org/10.1016/j.ergon.2011.05.001>.
- Avery, Eric, Troy Kelley, and Darush Davani. 2006. “Using Cognitive Architectures to Improve Robot Control: Integrating Production Systems, Semantic Networks, and Sub-Symbolic Processing.” In *Simulation Interoperability Standards Organization : 15th Conference on Behavior Representation in Modeling and Simulation*, 190–198.
- Baddeley, Alan. 2000. “The Episodic Buffer: A New Component of Working Memory?” *Trends in Cognitive Sciences* 4 (11): 417–423. [https://doi.org/10.1016/S1364-6613\(00\)01538-2](https://doi.org/10.1016/S1364-6613(00)01538-2).
- Baddeley, Alan. 2012. “Working Memory: Theories, Models, and Controversies.” *Annual Review of Psychology* 63 (1): 1–29. <https://doi.org/10.1146/annurev-psych-120710-100422>.
- Banks, Victoria A., Katherine L. Plant, and Neville A. Stanton. 2018. “Driver Error or Designer Error: Using the Perceptual Cycle Model to Explore the Circumstances Surrounding the Fatal Tesla Crash on 7th May 2016.” *Safety Science* 108:278–285. <https://doi.org/10.1016/j.ssci.2017.12.023>.
- Baron-Cohen, Simon, Alan M. Leslie, and Uta Frith. 1985. “Does the Autistic Child Have a ‘Theory of Mind’?” *Cognition* 21 (1): 37–46.
- Baud-Bovy, Gabriel, Pietro Morasso, Francesco Nori, Giulio Sandini, and Alessandra Sciutti. 2014. “Human Machine Interaction and Communication in Cooperative Actions.” In *Bioinspired Approaches for Human-Centric Technologies*, 241–268. Dordrecht: Springer. https://doi.org/10.1007/978-3-319-04924-3_8.

- Bauer, Andrea, Dirk Wollherr, and Martin Buss. 2008. "Human-Robot Collaboration: A Survey." *International Journal of Humanoid Robotics* 5 (1): 47–66. <https://doi.org/10.1142/S0219843608001303>.
- Bellet, Thierry, Béatrice Bailly-Asuni, Pierre Mayenobe, and Aurélie Banet. 2009. "A Theoretical and Methodological Framework for Studying and Modelling Drivers' Mental Representations." *Safety Science* 47 (9): 1205–1221. <https://doi.org/10.1016/j.ssci.2009.03.014>.
- Bemporad, Alberto, and Stefano Di Cairano. 2005. "Optimal Control of Discrete Hybrid Stochastic Automata." In *Lecture Notes in Computer Science*, 151–167. 3414. Berlin: Springer. https://doi.org/10.1007/978-3-540-31954-2_10.
- Berkes, Pietro, Gergo Orbán, Máté Lengyel, and József Fiser. 2011. "Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment." *Science* 331 (6013): 83–87. <https://doi.org/10.1126/science.1195870>.
- Breazeal, Cynthia. 2004. "Social Interactions in HRI: The Robot View." *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 34 (2): 181–186. <https://doi.org/10.1109/TSMCC.2004.826268>.
- Chen, Po Da, Chao Yin Wu, Rey Heng Hu, Chiung Nien Chen, Ray Hwang Yuan, Jin Tung Liang, Hong Shiee Lai, and Yao Ming Wu. 2017. "Robotic Major Hepatectomy: Is There a Learning Curve?" *Surgery (United States)* 161 (3): 642–649. <https://doi.org/10.1016/j.surg.2016.09.025>.
- Corredor, Javier, Jorge Sofrony, and Angelika Peer. 2017. "Decision-Making Model for Adaptive Impedance Control of Teleoperation Systems." *IEEE Transactions on Haptics* 10 (1): 5–16. <https://doi.org/10.1109/TOH.2016.2581807>.
- Curioni, Arianna, Gunther Knoblich, and Natalie Sebanz. 2017. "Joint Action in Humans: A Model for Human-Robot Interactions." In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 1–19. Dordrecht, Switzerland: Springer. https://doi.org/10.1007/978-94-007-7194-9_126-1.
- de Oliveira Filho, Getúlio Rodrigues. 2002. "The Construction of Learning Curves for Basic Skills in Anesthetic Procedures: An Application for the Cumulative Sum Method." *Anesthesia and Analgesia* 95 (2): 411–416. <https://doi.org/10.1213/0000539-200208000-00033>.
- de Winter, Joost C. F., Riender Happee, Marieke H. Martens, and Neville A. Stanton. 2014. "Effects of Adaptive Cruise Control and Highly Automated Driving on Workload and Situation Awareness: A Review of the Empirical Evidence." *Transportation Research Part F: Traffic Psychology and Behavior* 27:196–217. <https://doi.org/10.1016/j.trf.2014.06.016>.
- Di Cairano, Stefano, Daniele Bernardini, Alberto Bemporad, and Ilya V. Kolmanovskiy. 2014. "Stochastic MPC with Learning for Driver-Predictive Vehicle Control and Its Application to HEV Energy Management." *IEEE Transactions on Control Systems Technology* 22 (3): 1018–1031. <https://doi.org/10.1109/tcst.2013.2272179>.
- Dodd, Will, and Ridelto Gutierrez. 2005. "The Role of Episodic Memory and Emotion in a Cognitive Robot." In *IEEE International Workshop on Robot and Human Interactive Communication, 2005*, 692–697. New York: IEEE. <https://doi.org/10.1109/ROMAN.2005.1513860>.
- Dondrup, Christian, Nicola Bellotto, Marc Hanheide, Kerstin Eder, and Ute Leonards. 2015. "A Computational Model of Human-Robot Spatial Interactions Based on a Qualitative Trajectory Calculus." *Robotics* 4 (1): 63–102. <https://doi.org/10.3390/robotics4010063>.
- Doumerc, Nicolas, Carlo Yuen, Richard Savdie, M. Bayzidur Rahman, Kris K. Rasiah, Ruth Pe Benito, Warick Delprado, Jayne Matthews, Anne Maree Haynes, and Phillip D. Stricker. 2010. "Should Experienced Open Prostatic Surgeons Convert to Robotic Surgery? The Real Learning Curve for One Surgeon over 3 Years." *BJU International* 106 (3): 378–384. <https://doi.org/10.1111/j.1464-410X.2009.09158.x>.
- Eder, Kerstin, Chris Harper, and Ute Leonards. 2014. "Towards the Safety of Human-in-the-Loop Robotics: Challenges and Opportunities for Safety Assurance of Robotic Co-workers." In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 660–665. New York: IEEE. <https://doi.org/10.1109/ROMAN.2014.6926328>.
- Fagnant, Daniel J., and Kara Kockelman. 2015. "Preparing a Nation for Autonomous Vehicles: Opportunities, Barriers and Policy Recommendations." *Transportation Research Part A: Policy and Practice* 77:167–181. <https://doi.org/10.1016/j.tra.2015.04.003>.
- Fatemi, Mehdi, and Simon Haykin. 2014. "Cognitive Control: Theory and Application." *IEEE Access* 2:698–710. <https://doi.org/10.1109/ACCESS.2014.2332333>.
- Fischer, Tobias, and Yiannis Demiris. 2019. "Computational Modelling of Embodied Visual Perspective-Taking." *IEEE Transactions on Cognitive and Developmental Systems* 12 (4): 723–732. <https://doi.org/10.1109/TCDS.2019.2949861>.
- Fiser, József, Pietro Berkes, Gergo Orbán, and Máté Lengyel. 2010. "Statistically Optimal Perception and Learning: From Behavior to Neural Representations." *Trends in Cognitive Sciences* 14 (3): 119–130. <https://doi.org/10.1016/j.tics.2010.01.003>.
- Gold, Joshua I., and Hauke R. Heekeren. 2013. "Neural Mechanisms for Perceptual Decision Making." In *Neuroeconomics: Decision Making and the Brain*, edited by P. Glimcher and E. Fehr, 355–372. 2nd ed. San Diego: Elsevier. <https://doi.org/10.1016/B978-0-12-416008-8.00019-X>.

- Grigore, Elena Corina, Kerstin Eder, Anthony G. Pipe, Chris Melhuish, and Ute Leonards. 2013. "Joint Action Understanding Improves Robot-to-Human Object Handover." In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4622–4629. New York: IEEE. <https://doi.org/10.1109/IROS.2013.6697021>.
- Haefner, Ralf M., Pietro Berkes, and József Fiser. 2016. "Perceptual Decision-Making as Probabilistic Inference by Neural Sampling." *Neuron* 90 (3): 649–660. <https://doi.org/10.1016/j.neuron.2016.03.020>.
- Haykin, Simon, Mehdi Fatemi, Peyman Setoodeh, and Yanbo Xue. 2012. "Cognitive Control." *Proceedings of the IEEE* 100 (12): 3156–3169. <https://doi.org/10.1109/jproc.2012.2215773>.
- Herrmann, Guido, and Ute Leonards. 2018. "Human-Humanoid Interaction: Overview." In *Humanoid Robotics: A Reference*, edited by A. Goswami and P. Vadakkepat, 1–16. Dordrecht: Springer. https://doi.org/10.1007/978-94-007-7194-9_146-1.
- Hopper, A. N., M. H. Jamison, and W. G. Lewis. 2007. "Learning Curves in Surgical Practice." *Postgraduate Medical Journal* 83 (986): 777–779. <https://doi.org/10.1136/pgmj.2007.057190>.
- Houlsby, Neil M. T., Ferenc Huszár, Mohammad M. Ghassemi, Gergő Orbán, Daniel M. Wolpert, and Máté Lengyel. 2013. "Cognitive Tomography Reveals Complex, Task-Independent Mental Representations." *Current Biology* 23 (21): 2169–2175. <https://doi.org/10.1016/j.cub.2013.09.012>.
- Kaul, Sanjeev, Nikhil L. Shah, and Mani Menon. 2006. "Learning Curve Using Robotic Surgery." *Current Urology Reports*. Berlin: Springer. <https://doi.org/10.1007/s11934-006-0071-4>.
- Kawamura, Kazuhiko, and Stephen M. Gordon. 2006. "From Intelligent Control to Cognitive Control." In *2006 World Automation Congress, WAC'06*. New York: IEEE. <https://doi.org/10.1109/wac.2006.376003>.
- Kawamura, Kazuhiko, Stephen M. Gordon, Palis Ratanaswasd, Erdem Erdemir, and Joseph F. Hall. 2008. "Implementation of Cognitive Control for a Humanoid Robot." *International Journal of Humanoid Robotics* 5 (4): 547–586. <https://doi.org/10.1142/S0219843608001558>.
- Kellogg, Ronald Thomas. 2015. *Fundamentals of Cognitive Psychology*. 3rd ed. Thousand Oaks, CA: Sage.
- Khamassi, Mehdi, Stéphane Lallée, Pierre Enel, Emmanuel Procyk, and Peter F. Dominey. 2011. "Robot Cognitive Control with a Neurophysiologically Inspired Reinforcement Learning Model." *Frontiers in Neurobotics* 5:1. <https://doi.org/10.3389/fnbot.2011.00001>.
- Khan, Said G., Guido Herrmann, Frank L. Lewis, Tony Pipe, and Chris Melhuish. 2012. "Reinforcement Learning and Optimal Adaptive Control: An Overview and Implementation Examples." *Annual Reviews in Control* 36 (1): 42–59. <https://doi.org/10.1016/j.arcontrol.2012.03.004>.
- King, Dorothy, Angela Rowe, and Ute Leonards. 2011. "I Trust You; Hence I like the Things You Look At: Gaze Cueing and Sender Trustworthiness Influence Object Evaluation." *Social Cognition* 29 (4): 476–485. <https://doi.org/10.1521/soco.2011.29.4.476>.
- Kotseruba, Iuliia, and John K. Tsotsos. 2020. "40 Years of Cognitive Architectures: Core Cognitive Abilities and Practical Applications." *Artificial Intelligence Review* 53 (1): 17–94. <https://doi.org/10.1007/s10462-018-9646-y>.
- LaValle, Steven M. 2006. *Planning Algorithms*. Cambridge: Cambridge University Press.
- Levesque, Hector, and Gerhard Lakemeyer. 2008. "Cognitive Robotics." In *Foundations of Artificial Intelligence*, chap. 23. San Diego: Elsevier. [https://doi.org/10.1016/S1574-6526\(07\)03023-4](https://doi.org/10.1016/S1574-6526(07)03023-4).
- Lewis, Frank L., Dragana Vrabie, and Kyriakos G. Vamvoudakis. 2012. "Reinforcement Learning and Feedback Control: Using Natural Decision Methods to Design Optimal Adaptive Controllers." *IEEE Control Systems* 32 (6): 76–105. <https://doi.org/10.1109/mcs.2012.2214134>.
- Li, Zhijun, Yuanqing Xia, and Chun Yi Su. 2015. *Intelligent Networked Teleoperation Control*. Berlin: Springer. <https://doi.org/10.1007/978-3-662-46898-2>.
- Lopez Pulgarin, Erwin Jose, Guido Herrmann, and Ute Leonards. 2017. "Drivers' Manoeuvre Classification for Safe HRI." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 475–483. 10454 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-319-64107-2_37.
- Lopez Pulgarin, Erwin Jose, Guido Herrmann, and Ute Leonards. 2018. "Drivers' Manoeuvre Prediction for Safe HRI." In *IEEE International Conference on Intelligent Robots and Systems*, 8609–8614. New York: IEEE. <https://doi.org/10.1109/iros.2018.8593957>.
- Lopez Pulgarin, Erwin Jose, Tugrul Irmak, Joel Variath Paul, Arisara Meekul, Guido Herrmann, and Ute Leonards. 2018. "Comparing Model-Based and Data-Driven Controllers for an Autonomous Vehicle Task." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 170–182. 10965 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-319-96728-8_15.
- Maciejowski, J. M. 2002. *Predictive Control: With Constraints*. Essex, UK: Pearson Education.
- Mombaur, Katja, Anh Truong, and Jean Paul Laumond. 2010. "From Human to Humanoid Locomotion—an Inverse Optimal Control Approach." *Autonomous Robots* 28 (3): 369–383. <https://doi.org/10.1007/s10514-009-9170-7>.

- Morari, Manfred, and Jay H. Lee. 1999. "Model Predictive Control: Past, Present and Future." *Computers and Chemical Engineering* 23:667–682. [https://doi.org/10.1016/S0098-1354\(98\)00301-9](https://doi.org/10.1016/S0098-1354(98)00301-9).
- Moulin-Frier, Clement, Tobias Fischer, Maxime Petit, Gregoire Pointeau, Jordi Ysard Puigbo, Ugo Pattacini, Sock Ching Low, et al. 2018. "DAC-H3: A Proactive Robot Cognitive Architecture to Acquire and Express Knowledge about the World and the Self." *IEEE Transactions on Cognitive and Developmental Systems* 10 (4): 1005–1022. <https://doi.org/10.1109/tcds.2017.2754143>.
- Na, Jing, Muhammad Nasiruddin Mahyuddin, Guido Herrmann, Xuemei Ren, and Phil Barber. 2015. "Robust Adaptive Finite-Time Parameter Estimation and Control for Robotic Systems." *International Journal of Robust and Nonlinear Control* 25 (16): 3045–3071. <https://doi.org/10.1002/rnc.3247>.
- Na, Jing, Xuemei Ren, Cong Shang, and Yu Guo. 2012. "Adaptive Neural Network Predictive Control for Nonlinear Pure Feedback Systems with Input Delay." *Journal of Process Control* 22:194–206. <https://doi.org/10.1016/j.jprocont.2011.09.003>.
- Nakamura, Tomoaki, Takayuki Nagai, and Tadahiro Taniguchi. 2018. "SERKET: An Architecture for Connecting Stochastic Models to Realize a Large-Scale Cognitive Model." *Frontiers in Neurorobotics* 12 (12): 25. <https://doi.org/10.3389/fnbot.2018.00025>.
- Neerinx, Mark A., Willeke van Vught, Olivier Blanson Henkemans, Elettra Oleari, Joost Broekens, Rifca Peters, Frank Kaptein, et al. 2019. "Socio-cognitive Engineering of a Robotic Partner for Child's Diabetes Self-Management." *Frontiers in Robotics and AI* 6:118. <https://doi.org/10.3389/frobt.2019.00118>.
- Ogata, Katsuhiko. 2010. *Modern Control Engineering*. 5th ed. London: Pearson. <https://doi.org/10.1201/9781315214573>.
- Oh, Kwang-Myung, Ji-Hoon Kim, and Myung-Suk Kim. 2005. "Development of Humanoid Robot Design Process-Focused on the Concurrent Engineering Based Humanoid Robot Design." In *IDC International Design Congress 2005*, 1–13. International Design Congress. Yunlin, Taiwan: National Yunlin University of Science and Technology.
- Parker, Chris A. C., and Elizabeth A. Croft. 2012. "Design and Personalization of a Cooperative Carrying Robot Controller." In *Proceedings—IEEE International Conference on Robotics and Automation*, 3916–3921. New York: IEEE. <https://doi.org/10.1109/icra.2012.6225120>.
- Ratanaswasd, Palis, Stephen Gordon, and Will Dodd. 2005. "Cognitive Control for Robot Task Execution." In *IEEE International Workshop on Robot and Human Interactive Communication, 2005*, 440–445. New York: IEEE. <https://doi.org/10.1109/roman.2005.1513818>.
- Rosolia, Ugo, Xiaojing Zhang, and Francesco Borrelli. 2018. "Data-Driven Predictive Control for Autonomous Systems." *Annual Review of Control, Robotics, and Autonomous Systems* 1 (1): 259–286. <https://doi.org/10.1146/annurev-control-060117-105215>.
- SAE International. 2016. *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. https://doi.org/10.4271/J3016_201609.
- Samadi, David, Adam Levinson, Ari Hakimi, Ridwan Shabsigh, and Mitchell C. Benson. 2007. "From Proficiency to Expert, When Does the Learning Curve for Robotic-Assisted Prostatectomies Plateau? The Columbia University Experience." *World Journal of Urology* 25 (1): 105–110. <https://doi.org/10.1007/s00345-006-0137-4>.
- Scassellati, Brian. 2002. "Theory of Mind for a Humanoid Robot." *Autonomous Robots* 12 (1): 13–24. <https://doi.org/10.1023/A:1013298507114>.
- Sheng, Weihua, Anand Thobbi, and Ye Gu. 2015. "An Integrated Framework for Human-Robot Collaborative Manipulation." *IEEE Transactions on Cybernetics* 45 (10): 2030–2041. <https://doi.org/10.1109/tcyb.2014.2363664>.
- Spiers, Adam, Said Ghani Khan, and Guido Herrmann. 2016. *Biologically Inspired Control of Humanoid Robot Arms*. Cham, Switzerland: Springer.
- Tan, Huan, and Chen Liang. 2011. "A Conceptual Cognitive Architecture for Robots to Learn Behaviors from Demonstrations in Robotic Aid Area." In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1249–1252. New York: IEEE. <https://doi.org/10.1109/IEMBS.2011.6090294>.
- Turnbull, Oliver, Jonathan Lawry, Mark Lowenberg, and Arthur Richards. 2016. "A Cloned Linguistic Decision Tree Controller for Real-Time Path Planning in Hostile Environments." *Fuzzy Sets and Systems* 293:1–29. <https://doi.org/10.1016/j.fss.2015.08.017>.
- Visioli, Antonio, and Giovanni Legnani. 2002. "On the Trajectory Tracking Control of Industrial SCARA Robot Manipulators." *IEEE Transactions on Industrial Electronics* 49 (1): 224–232. <https://doi.org/10.1109/41.982266>.
- Warren, William. 2006. "The Dynamics of Perception and Action." *Psychological Review* 113 (2): 358–389. <http://search.proquest.com/docview/214221535/>.
- Wei, Changyun, and Koen V. Hindriks. 2013. "An Agent-Based Cognitive Robot Architecture." In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in*

Bioinformatics), edited by M. Dastani, J. F. Hübner, and B. Logan, 54–71. 7837 LNAI. Berlin: Springer. https://doi.org/10.1007/978-3-642-38700-5_4.

Whitsell, Bryan, and Panagiotis Artemiadis. 2017. “Physical Human-Robot Interaction (PHRI) in 6 DOF with Asymmetric Cooperation.” *IEEE Access* 5:10834–10845. <https://doi.org/10.1109/ACCESS.2017.2708658>.

Yamaguchi, Tomohiro, Yusuke Kinugasa, Akio Shiomi, Sumito Sato, Yushi Yamakawa, Hiroyasu Kagawa, Hiroyuki Tomioka, and Keita Mori. 2015. “Learning Curve for Robotic-Assisted Surgery for Rectal Cancer: Use of the Cumulative Sum Method.” *Surgical Endoscopy* 29 (7): 1679–1685. <https://doi.org/10.1007/s00464-014-3855-5>.

Yang, Weiwei, Guido Herrmann, Mark Lowenberg, and Xiaoqian Chen. 2010. “Dynamic Gain Scheduled Control in a Multi-variable Control Framework.” In *Proceedings of the IEEE Conference on Decision and Control*, 7081–7086. New York: IEEE. <https://doi.org/10.1109/cdc.2010.5717054>.

18 Social Cognition

Yukie Nagai

18.1 Introduction

People live in social environments. They receive various social signals, including gaze, facial expressions, gestures, and speech presented by other individuals. People also send such signals to others, regardless of their intentions. Most people have strong tendencies to attribute social meaning to the behavior of others and try to infer communicative intentions from them. Such tendencies enable other individuals, especially young children, to be easily engaged in social interactions.

This chapter addresses the issue of “social cognition.” It refers to the abilities of recognizing and controlling the self in relation to others and the abilities of applying and perceiving social signals in interactions with others. Such abilities enable people to exchange information and knowledge and thus to acquire new skills from other individuals. Indeed, people who have difficulties in social cognition often face challenges in learning new tasks. An example is autism spectrum disorder (ASD), a type of neurodevelopmental disorder characterized by deficits in social communication. People with low-functioning ASD often show difficulties in acquiring higher cognitive skills such as language use and cooperation. Investigating diverse social abilities is necessary to better understand the roles and the mechanisms of social cognition.

How can robotics researchers endow robots with humanlike social cognition? A promising approach is to learn from human infant development. It has been suggested that infants acquire basic social abilities in the first few years of life (Bremner 1994; Johnson 1997). They are born with limited abilities and gradually acquire physical and cognitive skills through interactions with the physical and social environment. In particular, caregivers play important roles in facilitating infant development. Caregivers engage their infants in social interactions and try to infer communicative signals from infant behavior. Such scaffoldings by caregivers enable infants to learn how to behave in social environments. Computational approaches inspired by infant development can be used to build cognitive developmental mechanisms in robots as well as to uncover underlying mechanisms of infant development (Asada et al. 2001; Asada et al. 2009; Cangelosi and Schlesinger 2015).

This chapter is organized as follows: First, psychology and neuroscience studies on social cognition are introduced in section 18.2. When and how infants acquire social cognitive

abilities are explained. We focus on four cognitive abilities that appear in early infancy: self-other recognition, joint attention, intention reading, and altruistic behaviors. Substantial findings relative to these abilities have motivated robotics researchers to replicate them in robots. Sections 18.3 to 18.5 describe robotics models of these cognitive functions. Computational models based on neural networks, probabilistic models, and reinforcement-learning models are introduced as potential mechanisms for development. Section 18.6 then presents a new developmental theory based on predictive coding. While a lot of robotics research has targeted a specific function in social cognition, the predictive coding theory provides a unified principle that accounts for both temporal continuity and individual diversity in development. Finally, section 18.7 concludes this chapter by presenting future issues.

18.2 Psychology and Neuroscience Theories on Social Cognition

18.2.1 Self-Other Recognition

Recognizing the self is a fundamental ability for infants (Bertenthal and Fischer 1978; Rochat 2003). Infants must discriminate their bodies from their environment in order to control their bodies in an intentional manner. They also need to differentiate other individuals from the self and the environment. Other individuals are active and self-propelled entities, who have similar but different bodies and internal states from those of infants. Detecting similarities as well as differences between the self and others is crucial for establishing social interactions.

Psychologists have been investigating when and how infants come to recognize themselves. Meltzoff, Saby, and Marshall (2019) examined neural representations of the self-body in sixty-day-old infants. They measured brain activities when infants received tactile stimulation of various body parts. Their results revealed differentiated body representations, which overlap with the body maps in the adult brain. The quality of body awareness changes during development. Infants at a few months of age are aware of their body as physical entities (Moore et al. 2007). They detect the relationship between proprioceptive and visual information and recognize their bodies visually. At around eighteen months of age, infants begin to recognize themselves even in reflections such as mirrors (Brownell, Zerwas, and Ramani 2007). Infants at this stage can pass a mirror test (Amsterdam 1972), which is a behavioral signature of self-recognition. A short time later, infants become able to recognize their bodies in time and space (Moore et al. 2007). They finally learn to effectively control their bodies in order to affect the environment.

Rochat (2003) summarized these findings and proposed six levels of self-awareness that unfold during infancy: confusion, differentiation, situation, identification, permanence, and self-consciousness. Infants start with confused representations of the self and the environment and gradually learn to differentiate their bodies from the environment. Later, they begin to identify their bodies in multiple modalities and finally extend self-awareness to time and space.

Despite a number of findings about self-awareness, the development of self-other recognition has been less studied. How infants detect similarities as well as differences between the self and others is still an open question. Some neural and behavioral evidence

supports a hypothesis that the self and others remain undifferentiated during early infancy. Meltzoff et al. (2018) found neural responses in the infant brain that detect the equivalence between the self and others. They revealed that the primary somatosensory cortex of the infant brain was activated for both a felt touch (i.e., being touched on their bodies) and an observed touch (i.e., observing another person being touched). Neonatal imitations are behavioral signatures to support the hypothesis. Meltzoff and Moore (1977) found that newborn babies could imitate facial and manual movements presented by other individuals. The ability to detect the equivalence between different modalities (e.g., vision and proprioception) and thus between the self and others might be innate in infants (Meltzoff and Moore 1977). Mirror neurons and mirror neuron systems are known to be relevant neural mechanisms for imitation (Rizzolatti et al. 1996; Rizzolatti, Fogassi, and Gallese 2001; Iacoboni and Dapretto 2006). It has been found that the same brain areas are activated both when a person is executing an action and when they are observing the same action performed by other individuals.

Section 18.3 presents robotics models for the development of self-recognition and self-other recognition. How to detect similarities as well as differences between the self and others is discussed from a computational standpoint.

18.2.2 Joint Attention

Joint attention is a phenomenon in which two people attend to the same object (Scaife and Bruner 1975; Moore and Dunham 1995). While self-other recognition is a dyadic interaction, joint attention concerns a triadic interaction involving an object. This ability is considered a cornerstone for many social abilities, including imitation, theory of mind, and language use, because it enables infants to share experiences with and learn from other persons (Tomassello and Farrar 1986; Charman et al. 2000; Morales et al. 2000).

Butterworth and Jarrett (1991) closely examined when and how infants come to achieve joint attention. They found three stages of development appearing from the age of six to eighteen months. The first stage is called the ecological stage. Infants from six to nine months old detect the direction of another person's gaze but cannot precisely localize the target object. Salient properties of the object rather than the gaze cue guide the infants' attention. The second stage is called the geometric stage, in which infants start establishing joint attention. At age twelve months, infants learn to follow the gaze direction of another person and identify the object the person is looking at. This ability, however, is limited to an object within the field of the infant's first view. Only in the third stage, called the representational stage, does the ability become fully functional. Infants at age eighteen months establish joint attention regardless of the position of the object. They can turn around along the direction of another person's gaze even when the object is located outside the field of the infants' first view. This stage requires a mental representation of the environment.

Other studies have investigated the effects of different visual cues on joint attention. The turning of another person's head facilitates gaze following in younger infants (Moore, Angelopoulos, and Bennett 1997). Infants who do not yet spontaneously follow a person's static head orientation can learn to follow a dynamic head turn. The coordinated movement of the head and the eyes together enable joint attention in younger infants (Lempers 1979). Only the orientation information or only eye movement is not sufficient for them to

perform joint attention. Furthermore, communicative signals from a person have been found to be crucial for joint attention (Senju and Csibra 2008). Six-month-old infants follow the direction of another person's gaze only when the person establishes eye contact or produces infant-directed speech, which is characterized by a wide range of pitch variation, before the gaze shift.

Inspired by these findings, various robotics models for joint attention have been proposed. Section 18.4 describes how computational studies have contributed to a better understanding of the underlying developmental mechanisms.

18.2.3 Reading Intentions and Altruistic Behavior

The abilities of reading intention and altruistic behavior are thought to be acquired based on joint attention. Once infants are able to share their experiences with others, they realize that other persons have unobservable internal states, such as beliefs and intentions. Theory of mind refers to this ability (Premack and Woodruff 1978; Baron-Cohen 1995), which becomes a basis for higher social cognition.

Woodward and colleagues (Woodward 1998; Sommerville, Woodward, and Needham 2005; Gerson and Woodward 2014) investigated when and how infants come to understand other persons' intentions. Employing visual habituation paradigms, they examined whether infants recognized a change in the goal of an experimenter's reaching action. Their results demonstrated that six-month-old infants could already encode another person's actions as goal directed (Woodward 1998). They further revealed that the motor experiences of infants have a great impact on infants' abilities with regard to action perception (Sommerville, Woodward, and Needham 2005; Gerson and Woodward 2014). Three-month-old infants, who could not yet spontaneously reach for an object, exhibited the ability to recognize goal directedness in another person's reaching after experiencing the apprehending of an object with a sticky mitten. The importance of motor experiences in reading others' intentions was also found in another study (Kanakogi and Itakura 2011). The researchers revealed synchronous development of action production and action perception in four- to 10-month-old infants.

The ability to infer another person's intentions can lead to the development of altruistic behavior. Older infants can help others by completing the others' goal even if they do not receive any immediate benefits. Warneken and Tomasello (2006) showed that eighteen-month-old infants could help others in a variety of different situations, such as handing over an out-of-reach object, opening a cabinet to store objects, and so on. Younger infants, in contrast, could help other persons only in easier scenarios (Warneken and Tomasello 2007). Cirelli et al. (Cirelli, Einarson, and Trainor 2014; Cirelli, Wan, and Trainor 2014) examined prerequisites for altruism in infants. They revealed that the social relationship between infants and an experimenter affects helpfulness in infants. Fourteen-month-old infants more significantly helped an experimenter who presented body movement synchronized with infants versus asynchronous movement.

An open question is the developmental mechanism and the motivation for altruistic behaviors. Two hypotheses have been proposed to account for development (Paulus 2014). The first hypothesis is called the emotion-sharing model, which proposes that infants are able to differentiate the self from others and to recognize others as intentional agents.

Infants are motivated to help other persons based on empathic concerns for the needs of others. The second hypothesis is called the goal-alignment model. This model does not propose self-other discrimination but rather assumes undifferentiated self-other representations. It is believed that infants estimate the goal of other persons and take over the goal as if it were their own.

Section 18.5 presents robotics models for reading intentions and altruistic behavior. Computational studies provide new insight into how these abilities successively develop through a common mechanism.

18.3 Cognitive Robotics Models for Self-Other Recognition

This section presents computational models for self-recognition and self-other recognition. Refer to a recent review article (Georgie, Schillaci, and Hafner 2019) for more details.

18.3.1 Robotics Models for Self-Recognition

Many computational models have been proposed to enable robots to recognize their own bodies. Yamada et al. (2016) and Hoffmann et al. (2018) investigated the development of body representations in artificial agents. They proposed learning models for a fetus simulator or a humanoid robot to self-organize somatosensory signals through tactile experiences. Their experiments demonstrated structured body representations acquired in the simulator/robot that were similar to those found in the primary somatosensory cortex in humans or primates. Hafner and colleagues (Lang, Schillaci, and Hafner 2018; Schillaci, Hafner, and Lara 2016) and Lanillos and colleagues (Lanillos, Dean-Leon, and Cheng 2017; Lanillos and Cheng 2018) developed learning models for robots to visually recognize their own bodies. Their key idea was that a forward model that learns to predict sensory signals through multimodal experiences plays a crucial role in self-recognition. Their experiments replicated not only self-recognition ability but also relevant phenomena such as attenuations of self-generated movements (Lang, Schillaci, and Hafner 2018) and rubber hand illusions (Lanillos and Cheng 2018). In contrast to these studies focusing on sensory predictability, Tani (1998) suggested that the self becomes aware through interactions between the bottom-up sensations and the top-down predictions in dynamic systems. Neural networks, which are trained to achieve certain goals, transit spontaneously between goal-directed stable states and unstable states. His study on analogies between the model's behaviors and the literature on the phenomenology of self-recognition indicated that the self is recognized during unstable phases.

18.3.2 Robotics Models for Self-Other Recognition

Other researchers have addressed the issue of self-other recognition. Gold and Scassellati (2009) proposed a probabilistic model for a robot to discriminate its own body, the bodies of animate individuals, and inanimate objects (see figure 18.1a). They hypothesized that these entities could be detected as image motion that has different probabilities of being generated by the robot's motor commands. Their experiment demonstrated that the robot could successfully differentiate its body from the body of a human even using an image

reflected in a mirror. Nakajo et al. (2016) proposed a recurrent neural network that differentiates the self, other individuals, and objects based on the certainty of predictions (see figure 18.1b). Their network could learn to predict the variance as well as the mean of sensory signals, where the variance was used as an index of predictability. Their experiments demonstrated successful discrimination of the robot's own body as a highly predictable entity (i.e., low variance) compared to other persons or objects.

In contrast to the above studies focusing on self-other discrimination, Nagai et al. (Nagai, Kawai, and Asada 2011; Kawai, Nagai, and Asada 2012) proposed a neural network that learns to detect both similarities and differences between the self and others (see figure 18.1c). They hypothesized that immaturity in sensory acuity enhances self-other equivalence in the early stage of development and therefore enables the network to maintain the self-other correspondence while learning to differentiate the self and others. Their experiments comparing different learning conditions acknowledged the importance of sensory development. Only the network with sensory development acquired both similarities and differences between the self and others. All the above computational studies provide important insights into the underlying neural mechanisms for self-other recognition.

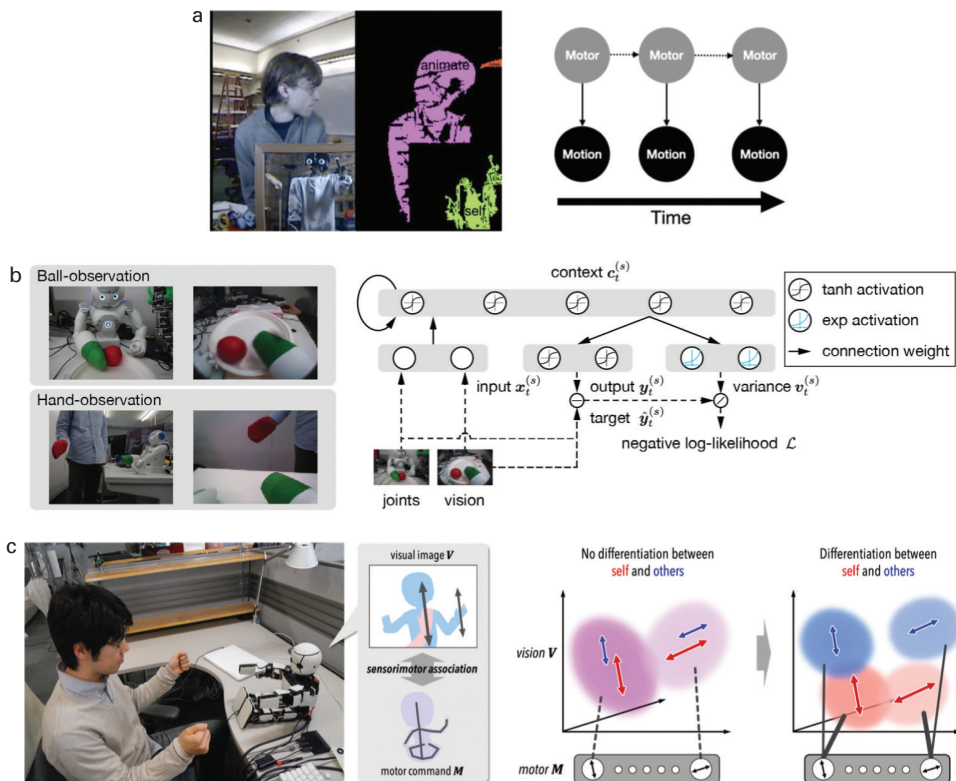


Figure 18.1

Cognitive robotics models for self-other recognition. (a) Self-other recognition using probabilistic Bayesian models (Gold and Scassellati 2009). (b) Self-other recognition using a recurrent neural network with variance prediction (Nakajo et al. 2016). (c) Self-other recognition based on sensorimotor associative learning with sensory development (Nagai, Kawai, and Asada 2011).

18.4 Cognitive Robotics Models for Joint Attention

Various robotics models for joint attention have been proposed, inspired by behavioral findings about infant development. This section presents computational models using different learning architectures, such as neural networks and reinforcement learning. Refer to Kaplan and Hafner (2006) for a comprehensive review of joint attention in robots.

18.4.1 Neural Network Models for Joint Attention

Studies by Nagai and colleagues (Nagai et al. 2003; Nagai, Asada, and Hosoda 2006; Nagai 2005b) proposed neural network models through which robots learned to achieve joint attention with human caregivers (see figure 18.2a). Their networks were designed to learn the sensorimotor contingency between a visual input (i.e., a camera image capturing the caregiver's face) and a motor output (i.e., a motor command to shift the robot's gaze direction). Their key ideas were that only successful experiences of joint attention involve higher sensorimotor correlations and that these correlations can be acquired by a network even without explicit teaching signals. Their experiments not only replicated behavioral findings from psychology but also provided new insights into the underlying mechanisms. For example, unsupervised contingency learning could lead to the three-staged development of joint attention, as observed in infants (refer to section 18.2.2; Nagai et al. 2003); sensory development and caregiver's scaffolding could facilitate learning (Nagai, Asada, and Hosoda 2006); and motion information from the caregiver's head turn could enable early development of joint attention, as observed in young infants (Nagai 2005b). Nagai (2005a) further applied a neural network to the development of comprehending deictic gestures. She trained a robot to recognize human gestures such as reaching, tapping, and pointing. The experiment demonstrated that reaching gestures were easier to recognize than the other two, as observed in infants. The static and motion cues produced by reaching gestures were richer and thus contributed to earlier development.

18.4.2 Reinforcement-Learning Models for Joint Attention

In contrast to the above studies using neural networks, studies by Triesch and colleagues (Jasso et al. 2012; Triesch, Jasso, and Deák 2007; Triesch et al. 2006) proposed joint attention models based on reinforcement learning (see figure 18.2b). Their key idea was that an infant learner acquires a sensorimotor map based on the rewards of looking at a salient object. The sensorimotor signals used in their experiment included a saliency map of the environment, the head and eye direction of a caregiver, and the gaze direction of the infant. Their experiments replicated multiple aspects of joint attention: the staged development of joint attention, facilitated learning with head and eye cues from a caregiver (Jasso et al. 2012), and mirror-neuron-like properties acquired in motor representations (Triesch, Jasso, and Deák 2007). They also examined the causes of developmental delays or difficulties observed in ASD and Williams syndrome. Their experiments manipulating model parameters suggested that atypical reward structures for the caregiver's face and objects prevented the development of joint attention (Triesch et al. 2006).

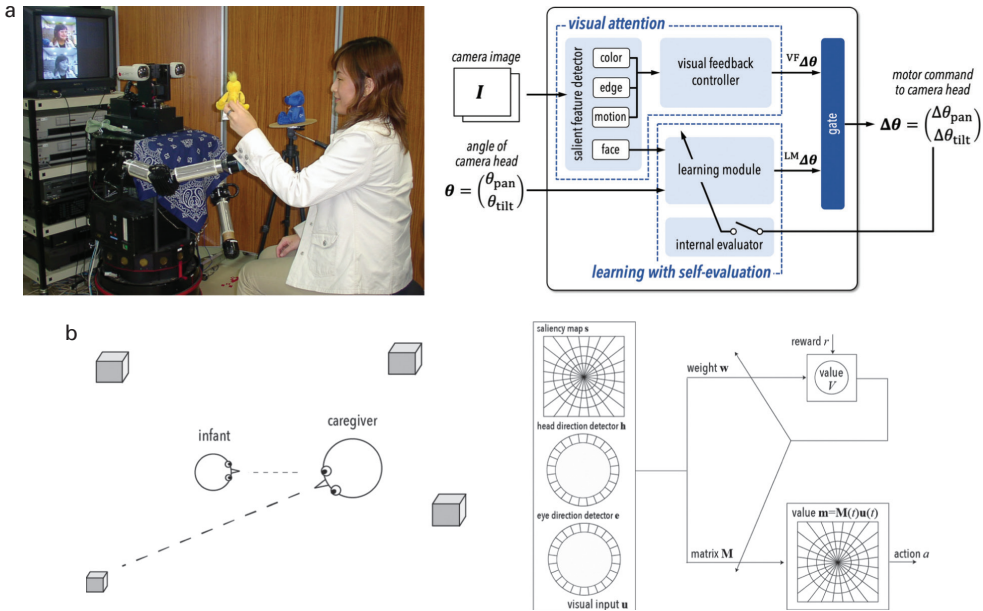


Figure 18.2

Cognitive robotics models for joint attention. (a) Development of joint attention based on contingency learning using a neural network (Nagai et al. 2003). (b) Development of joint attention using reinforcement learning (Jasso et al. 2012).

18.4.3 Miscellaneous Models

Sumioka, Yoshikawa, and Asada (2008, 2010) extended the idea of contingency learning proposed by Nagai et al. (2003). They assumed that robots as well as infants do not know what sensorimotor signals to learn beforehand. They employed transfer entropy to detect inherent contingency in social interactions. Their experiments demonstrated the successful open-ended development of joint attention and relevant functions (e.g., gaze following and gaze alternation). Hoffman et al. (2006) proposed a probabilistic model combined with a saliency map. Inspired by active intermodal mapping as a basis for infant imitation (Meltzoff and Moore 1997), their model learned supramodal representations between the visual and proprioceptive signals of a robot. Their experiments showed that learned probability distribution represented even instructor-specific distributions over objects.

18.5 Cognitive Robotics Models for Reading Intentions and Altruistic Behavior

How people infer the internal states (e.g., intentions and emotions) of other individuals remains unclear in social cognition. People cannot directly access other persons' internal states and do not always receive feedback from others. Unlike the abilities of self-other recognition and joint attention, statistical learning through sensorimotor experiences is not sufficient for the development of these abilities. This section presents robotics studies that have addressed this challenge.

18.5.1 Robotics Models for Reading Intentions

Inspired by the discovery of mirror neurons (Rizzolatti et al. 1996; Rizzolatti, Fogassi, and Gallese 2001; Iacoboni and Dapretto 2006) robotics researchers have proposed neural network models that exhibit activation similar to mirror neurons. Neuroscience studies have shown that the human brain recognizes the goal and estimates the intention of another individual's actions by recruiting the brain areas used for action generation. Copete, Nagai, and Asada (2016) replicated this neural function using a deep autoencoder (see figure 18.3a). The network was first trained through a robot's motor experience to execute desired actions (e.g., reaching) and was then applied to recognizing actions by other individuals. Of importance is that the robot received only the visual input during action observations, while it obtained the visual, tactile, and proprioceptive signals during action executions. Their key ideas were that the network could reconstruct unobservable signals through multimodal representations and that the reconstructed signals could be used for further prediction of future sensory states. Their experiments demonstrated that the imaginary tactile and proprioceptive signals recalled from the visual input contributed to a more accurate estimation of future states, which is indicative of intention reading.

Horii, Nagai, and Asada (2016, 2018) proposed a multimodal deep belief network able to estimate and imitate the emotions of others (see figure 18.3b). Emotional states such as happiness and sadness are internal states and must be inferred from observable signals. Their key idea was analogous to the model put forth by Copete, Nagai, and Asada (2016). Multimodal representations acquired through one's own motor experiences enable the network to estimate

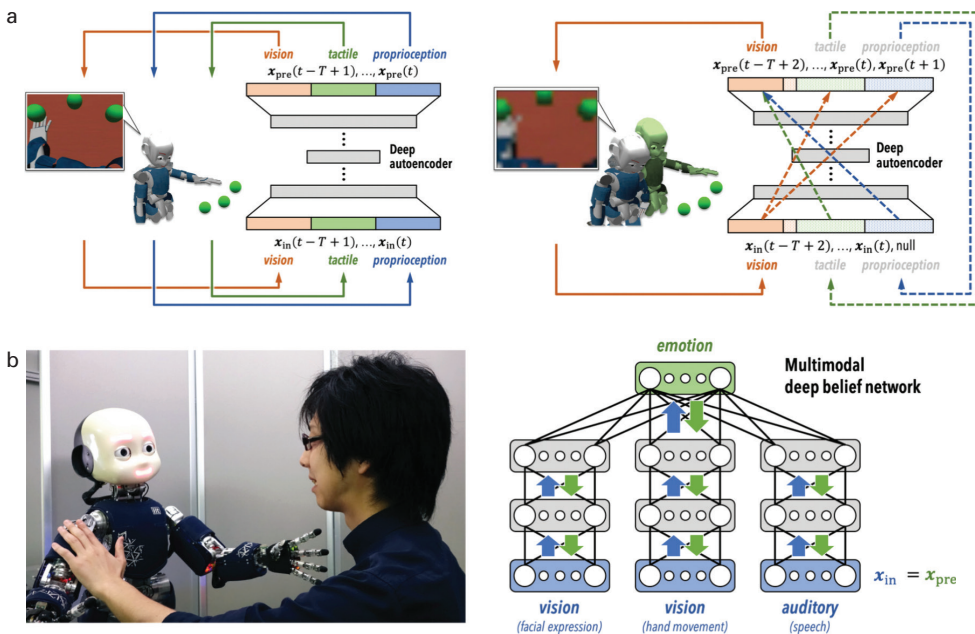


Figure 18.3 Cognitive robotics models for reading intention. (a) Reading intention based on mirror neurons using a deep autoencoder (Copete, Nagai, and Asada 2016). (b) Estimation and imitation of emotion using a multimodal deep belief network (Horii, Nagai, and Asada 2016).

the internal states and furthermore improve the estimation by reconstructing unobservable sensory signals. Their experiments demonstrated that a robot equipped with the network could acquire emotional states through developmental differentiation, as in infants (Horii, Nagai, and Asada 2018), and that it could estimate and imitate emotional states of humans (Horii, Nagai, and Asada 2016).

18.5.2 Robotics Models for Altruistic Behavior

Once robots are able to estimate the intention of other individuals, they can help others by completing others' goals. Baraglia et al. (Baraglia, Nagai, and Asada 2016; Baraglia et al. 2017) proposed a robotics model for altruistic behavior by extending the model proposed by Copete, Nagai, and Asada (2016). They suggested that a robot equipped with a mirror-neuron-like mechanism is able to estimate the goal of another person based on the robot's own motor experiences and furthermore is able to fulfill the goal as if it were their own. This idea supports the goal-alignment model hypothesis (Paulus 2014). The robot does not need to differentiate the self from others but rather exploits an immature representation between them. Their experiment replicated the developmental progress observed in infants (Warneken and Tomasello 2006, 2007). The robot that had more experiences of action generation produced helping behavior in wider situations, whereas the robot with less action experiences showed limited abilities. This result supports the developmental hypothesis and further provides a potential neural mechanism for altruism.

18.6 A Unified Computational Theory for Social Cognition

Most studies in cognitive robotics have focused on a specific ability of social cognition. For example, one model reproduced the ability of self-other recognition but did not address joint attention. In contrast, psychological studies suggest that cognitive abilities are closely interlinked. An open challenge for robotics researchers is to propose a unified mechanism that drives the continuous development of multiple cognitive functions.

18.6.1 Cognitive Development Based on Predictive Coding

Nagai (2019) suggests that the theory of predictive coding provides a unified account for both temporal continuity and individual diversity in cognitive development. The predictive coding theory was originally proposed as a principle of the human brain (Friston, Kilner, and Harrison 2006; Friston 2010; Clark 2013). Neuroscientists suggest that the brain works as a predictive machine that tries to minimize prediction errors between incoming sensory signals and top-down predictions produced by internal models. Of importance is that both perception and action are produced through the process of minimizing prediction errors. Perceptions are formed by integrating sensory signals with top-down predictions according to their precision (i.e., perceptual inference), whereas actions are generated to minimize prediction errors by altering sensory signals (i.e., active inference).

Nagai (2019) suggested that two processes of minimizing prediction errors lead to the continuous development of social cognition (see figure 18.4). First, the process of updating internal models enables infants to acquire basic sensorimotor abilities related to the self (see figure 18.4a). Humans are born with immature internal models and therefore must update their models through sensorimotor experiments. For example, the abilities of self-




Figure 18.4

Cognitive development based on predictive learning. (a) Updating the internal model through own sensorimotor experiences. (b) Generating actions to alter sensory signals.

recognition and self-other differentiation are achieved by detecting their different predictabilities using the internal models. Goal-directed actions are acquired by updating the internal models in order to intentionally control the self. Internal models come to represent the relationship between proprioceptive signals and exteroceptive/interoceptive signals (e.g., vision and audition) from the body.

Second, social cognitive abilities are considered to emerge through the process of acting on the environment to minimize prediction errors (see figure 18.4b). When interacting with other individuals, infants detect prediction errors because the behavior of others cannot be completely predicted using their internal models. Generating actions to minimize prediction errors results in protosocial behavior. For example, altruistic behavior emerges as a process of completing predicted goals, which were thought to be achieved by other persons. This view agrees with the goal-alignment model (Paulus 2014) and suggests that early forms of altruistic behavior do not involve social motivation. Only in the later stage of development do infants acquire social motivation by receiving social feedback.

18.6.2 ASD Caused by Impairments in Predictive Processing

The development of social cognition shows individual diversity. Some infants exhibit developmental delays and/or difficulties in acquiring cognitive functions. ASD is a type of neurodevelopmental disorder characterized by difficulties with social communication and interaction and a preference for restricted and repetitive patterns of behaviors, interests, and activities (American Psychiatric Association 2013). Despite substantial behavioral and neural evidence about ASD, its developmental cause has not been fully elucidated.

Inspired by the predictive coding theory (Friston 2010; Friston, Kilner, and Harrison 2006; Clark 2013), neuroscientists have suggested that impairments in predictive processing may produce the hypersensory sensitivities and/or lower adaptabilities observed in ASD (Pellucano and Burr 2012; Brock 2012; Van de Cruys et al. 2014). Diverse characteristics of ASD might be accounted for by too weak or too strong reliance on predictions (Nagai 2019). Two computational studies that have tested this hypothesis are presented below. Refer to Lanillos et al. (2020) for a comprehensive review.

Idei et al. (2018) investigated the influence of the precision of sensory predictions on a robot's behaviors (see figure 18.5a). They employed a recurrent neural network called

S-CTRNN (Murata et al. 2013) that can learn to predict not only sensory inputs but also their variances based on the minimization of precision-weighted prediction errors. The network implemented in a robot was first trained with ball manipulation tasks and then tested to reproduce the tasks using a modified model parameter. Their results demonstrated that both increased and decreased sensory precision induced behavioral rigidity similar to ASD. Decreased sensory precision caused invariability of the robot's intention, whereas increased sensory precision resulted in fluctuations and subsequent fixations of the intention.

Philippesen and Nagai (2018) investigated how hyper- and hypo-priors affect predictive learning (see figure 18.5*b*). They employed S-CTRNN (Murata et al. 2013), the same network used in Idei et al. (2018), and altered model parameters that control hyper- and hypo-priors in predictive processing. In contrast to the previous study, their experiment modified the parameters both during and after learning because properties of ASD should emerge during development. Their experiments demonstrated that ASD-like behaviors emerged with two extremes of model parameters, whereas behaviors similar to typically developed individuals were produced with properly balanced parameters. On one hand, hyper-priors prevented the network from learning to achieve the tasks because the network strongly relied on its own dynamics and ignored target signals. On the other hand, hypo-priors achieved very precise task behaviors but failed to acquire generalization capabilities. The internal representations of the network were unstructured because it did not utilize its own dynamics and was overfitted to the target behaviors. They concluded that a high variety of ASD behaviors could be accounted for by two extremes of hyper- and hypo-priors in predictive processing.

18.7 Conclusion

This chapter introduced social cognition from the perspective of psychology, neuroscience, and robotics. Human infants acquire social cognitive abilities such as self-other recognition, joint attention, intention reading, and altruistic behaviors through interactions with other individuals. A number of findings from psychology and neuroscience have motivated robotics researchers to design computational models for social cognition. Conversely, such models have contributed to a better understanding of the underlying mechanisms for cognitive development. Future issues to be addressed include a closer verification of the new theory of predictive coding that provides a unified account for cognitive development. To what extent the theory explains different aspects of development should be investigated from both analytical and synthetic approaches.

Additional Reading and Resources

- This paper presents the details about a developmental theory based on predictive coding (relevant to section 18.6). It explains to what extent the predictive coding theory can account for temporal continuity and individual diversity in cognitive development, with examples of robotic experiments: Nagai, Yukie. 2019. "Predictive Learning: Its Key Role in Early Cognitive Development." *Philosophical Transactions of the Royal Society B: Biological Sciences* 374 (1771): 20180030.

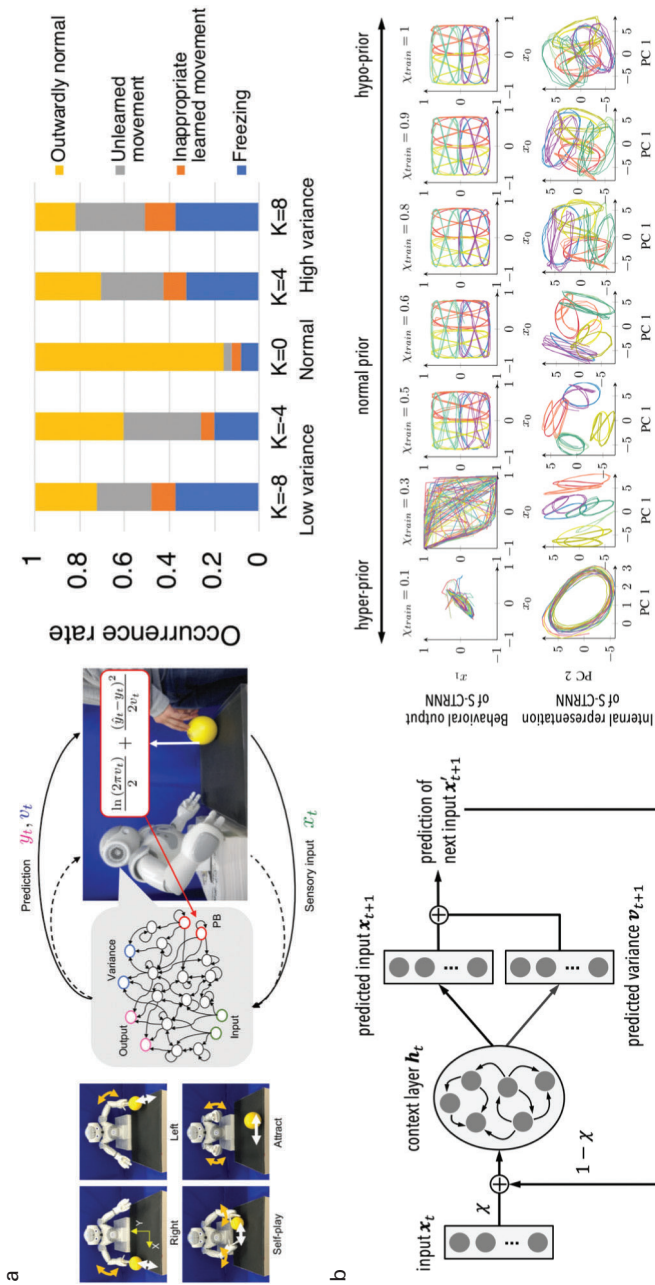


Figure 18.5
Cognitive robotics models for investigating underlying mechanisms for ASD. (a) ASD-like behaviors generated by increased/decreased sensory precisions (Idei et al. 2018). (b) ASD-like behaviors generated by hyper/hypo-priors in network learning (Philippssen and Nagat 2018).

- This paper provides a comprehensive review of computational studies on autism spectrum disorder and schizophrenia (relevant to section 18.6.2): Lanillos, Pablo, Daniel Oliva, Anja Philippsen, Yuichi Yamashita, Yukie Nagai, and Gordon Cheng. 2020. “A Review on Neural Network Models of Schizophrenia and Autism Spectrum Disorder.” *Neural Networks* 122:338–363.
- This paper presents a neural network model for the development of joint attention. The robot experiment demonstrates the three-staged development as observed in human infants (relevant to section 18.4.1): Nagai, Yukie, Koh Hosoda, Akio Morita, and Minoru Asada. 2003. “A Constructive Model for the Development of Joint Attention.” *Connection Science* 15 (4): 211–229.
- Project on cognitive and developmental robotics, including social cognition modeling: JSPS Grant-in-Aid for Specially Promoted Research “Constructive Developmental Science”: <https://www.youtube.com/watch?v=1etzhzSd171>.
- Project on cognitive robotics for cognitive modeling: JST CREST “Cognitive Mirroring” (in Japanese): <https://www.jst.go.jp/kisoken/jyonetsu/interview/h29/nagai.html>; <https://www.youtube.com/watch?v=1Onr2xssces>.

References

- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders (DSM-5®)*. Washington, DC: American Psychiatric.
- Amsterdam, Beulah. 1972. “Mirror Self-Image Reactions before Age Two.” *Developmental Psychobiology* 5 (4): 297–305.
- Asada, Minoru, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. 2009. “Cognitive Developmental Robotics: A Survey.” *IEEE Transactions on Autonomous Mental Development* 1 (1): 12–34.
- Asada, Minoru, Karl F. MacDorman, Hiroshi Ishiguro, and Yasuo Kuniyoshi. 2001. “Cognitive Developmental Robotics as a New Paradigm for the Design of Humanoid Robots.” *Robotics and Autonomous Systems* 37 (2–3): 185–193.
- Baldwin, D. A., C. Moore, and P. J. Dunham. 1995. “Joint Attention: Its Origins and Role in Development.” *Understanding the Link between Joint Attention and Language* 131:158.
- Baraglia, Jimmy, Maya Cakmak, Yukie Nagai, Rajesh P. N. Rao, and Minoru Asada. 2017. “Efficient Human-Robot Collaboration: When Should a Robot Take Initiative?” *International Journal of Robotics Research* 36 (5–7): 563–579.
- Baraglia, Jimmy, Jorge L. Copete, Yukie Nagai, and Minoru Asada. 2015. “Motor Experience Alters Action Perception through Predictive Learning of Sensorimotor Information.” In *2015 Joint IEEE 5th International Conference on Development and Learning and Epigenetic Robotics*, 63–69. New York: IEEE.
- Baraglia, Jimmy, Yukie Nagai, and Minoru Asada. 2016. “Emergence of Altruistic Behavior through the Minimization of Prediction Error.” *IEEE Transactions on Cognitive and Developmental Systems* 8 (3): 141–151.
- Baron-Cohen, S. 1995. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press.
- Bertenthal, Bennett I., and Kurt W. Fischer. 1978. “Development of Self-Recognition in the Infant.” *Developmental Psychology* 14 (1): 44.
- Bremner, J. Gavin. 1994. *Infancy*. Hoboken, NJ: Blackwell.
- Brock, Jon. 2012. “Alternative Bayesian Accounts of Autistic Perception: Comment on Pellicano and Burr.” *Trends in Cognitive Sciences* 16 (12): 573.
- Brownell, Celia A., Stephanie Zerwas, and Geetha B. Ramani. 2007. “‘So Big’: The Development of Body Self-Awareness in Toddlers.” *Child Development* 78 (5): 1426–1440.
- Butterworth, George, and Nicholas Jarrett. 1991. “What Minds Have in Common is Space: Spatial Mechanisms Serving Joint Visual Attention in Infancy.” *British Journal of Developmental Psychology* 9 (1): 55–72.

- Cangelosi, Angelo, and Matthew Schlesinger. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Charman, Tony. 2003. "Why Is Joint Attention a Pivotal Skill in Autism?" *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 358 (1430): 315–324.
- Charman, Tony, Simon Baron-Cohen, John Swettenham, Gillian Baird, Antony Cox, and Auriol Drew. 2000. "Testing Joint Attention, Imitation, and Play as Infancy Precursors to Language and Theory of Mind." *Cognitive Development* 15 (4): 481–498.
- Cirelli, Laura K., Kathleen M. Einarson, and Laurel J. Trainor. 2014. "Interpersonal Synchrony Increases Prosocial Behavior in Infants." *Developmental Science* 17 (6): 1003–1011.
- Cirelli, Laura K., Stephanie J. Wan, and Laurel J. Trainor. 2014. "Fourteen-Month-Old Infants Use Interpersonal Synchrony as a Cue to Direct Helpfulness." *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1658): 20130400.
- Clark, Andy. 2013. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences* 36 (3): 181–204.
- Copete, Jorge Luis, Yukie Nagai, and Minoru Asada. 2016. "Motor Development Facilitates the Prediction of Others' Actions through Sensorimotor Predictive Learning." In *2016 Joint IEEE 6th International Conference on Development and Learning and Epigenetic Robotics*, 223–229. New York: IEEE.
- Friston, Karl. 2010. "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11 (2): 127–138.
- Friston, Karl, James Kilner, and Lee Harrison. 2006. "A Free Energy Principle for the Brain." *Journal of Physiology-Paris* 100 (1–3): 70–87.
- Gallagher, Shaun. 2000. "Philosophical Conceptions of the Self: Implications for Cognitive Science." *Trends in Cognitive Sciences* 4 (1): 14–21.
- Georgie, Yasmin Kim, Guido Schillaci, and Verena Vanessa Hafner. 2019. "An Interdisciplinary Overview of Developmental Indices and Behavioral Measures of the Minimal Self." In *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics*, 129–136. New York: IEEE.
- Gerson, Sarah A., and Amanda L. Woodward. 2014. "Learning from Their Own Actions: The Unique Effect of Producing Actions on Infants' Action Understanding." *Child Development* 85 (1): 264–277.
- Gold, Kevin, and Brian Scassellati. 2009. "Using Probabilistic Reasoning over Time to Self-Recognize." *Robotics and Autonomous Systems* 57 (4): 384–392.
- Hoffman, Matthew W., David B. Grimes, Aaron P. Shon, and Rajesh P. N. Rao. 2006. "A Probabilistic Model of Gaze Imitation and Shared Attention." *Neural Networks* 19 (3): 299–310.
- Hoffmann, Matej, Zdeněk Straka, Igor Farkaš, Michal Vavrečka, and Giorgio Metta. 2018. "Robotic Homunculus: Learning of Artificial Skin Representation in a Humanoid Robot Motivated by Primary Somatosensory Cortex." *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 163–176.
- Horii, Takato, Yukie Nagai, and Minoru Asada. 2016. "Imitation of Human Expressions Based on Emotion Estimation by Mental Simulation." *Paladyn, Journal of Behavioral Robotics* 7 (1):40–54.
- Horii, Takato, Yukie Nagai, and Minoru Asada. 2018. "Modeling Development of Multimodal Emotion Perception Guided by Tactile Dominance and Perceptual Improvement." *IEEE Transactions on Cognitive and Developmental Systems* 10 (3): 762–775.
- Iacoboni, Marco, and Mirella Dapretto. 2006. "The Mirror Neuron System and the Consequences of Its Dysfunction." *Nature Reviews Neuroscience* 7 (12): 942–951.
- Idei, Hayato, Shingo Murata, Yiwen Chen, Yuichi Yamashita, Jun Tani, and Tetsuya Ogata. 2018. "A Neurorobotics Simulation of Autistic Behavior Induced by Unusual Sensory Precision." *Computational Psychiatry* 2:164–182.
- Jasso, Hector, Jochen Triesch, Gedeon Deák, and Joshua M. Lewis. 2012. "A Unified Account of Gaze Following." *IEEE Transactions on Autonomous Mental Development* 4 (4): 257–272.
- Johnson, Mark H. 1997. *Developmental Cognitive Neuroscience*. Hoboken, NJ: Blackwell.
- Kanakogi, Yasuhiro, and Shoji Itakura. 2011. "Developmental Correspondence between Action Prediction and Motor Ability in Early Infancy." *Nature Communications* 2 (1): 1–6.
- Kaplan, Frederic, and Verena V. Hafner. 2006. "The Challenges of Joint Attention." *Interaction Studies* 7 (2): 135–169.
- Kawai, Yuji, Yukie Nagai, and Minoru Asada. 2012. "Perceptual Development Triggered by Its Self-Organization in Cognitive Learning." In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5159–5164. New York: IEEE.
- Lang, Claus, Guido Schillaci, and Verena V. Hafner. 2018. "A Deep Convolutional Neural Network Model for Sense of Agency and Object Permanence in Robots." In *2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics*, 257–262. New York: IEEE.

- Lanillos, Pablo, and Gordon Cheng. 2018. "Adaptive Robot Body Learning and Estimation through Predictive Coding." In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4083–4090. New York: IEEE.
- Lanillos, Pablo, Emmanuel Dean-Leon, and Gordon Cheng. 2017. "Yielding Self-Perception in Robots through Sensorimotor Contingencies." *IEEE Transactions on Cognitive and Developmental Systems* 9 (2): 100–112.
- Lanillos, Pablo, Daniel Oliva, Anja Philippsen, Yuichi Yamashita, Yukie Nagai, and Gordon Cheng. 2020. "A Review on Neural Network Models of Schizophrenia and Autism Spectrum Disorder." *Neural Networks* 122:338–363.
- Lempers, Jacques D. 1979. "Young Children's Production and Comprehension of Nonverbal Deictic Behaviors." *Journal of Genetic Psychology* 135 (1): 93–102.
- Lombardo, Michael V., Bhismadev Chakrabarti, Edward T. Bullmore, Susan A. Sadek, Greg Pasco, Sally J. Wheelwright, John Suckling, MRC Aims Consortium, and Simon Baron-Cohen. 2010. "Atypical Neural Self-Representation in Autism." *Brain* 133 (2): 611–624.
- Meltzoff, Andrew N., and M. Keith Moore. 1977. "Imitation of Facial and Manual Gestures by Human Neonates." *Science* 198 (4312): 75–78.
- Meltzoff, Andrew N., and M. Keith Moore. 1997. "Explaining Facial Imitation: A Theoretical Model." *Infant and Child Development* 6 (3–4): 179–192.
- Meltzoff, Andrew N., Rey R. Ramirez, Joni N. Saby, Eric Larson, Samu Taulu, and Peter J. Marshall. 2018. "Infant Brain Responses to Felt and Observed Touch of Hands and Feet: An MEG Study." *Developmental Science* 21 (5): e12651.
- Meltzoff, Andrew N., Joni N. Saby, and Peter J. Marshall. 2019. "Neural Representations of the Body in 60-Day-Old Human Infants." *Developmental Science* 22 (1): e12698.
- Moore, C., and P. J. Dunham. 1995. *Joint Attention: Its Origins and Role in Development*. Mahway, NJ: Lawrence Erlbaum.
- Moore, Chris, Maria Angelopoulos, and Paula Bennett. 1997. "The Role of Movement in the Development of Joint Visual Attention." *Infant Behavior and Development* 20 (1): 83–92.
- Moore, Chris, Jennifer Mealiea, Nancy Garon, and Daniel J. Povinelli. 2007. "The Development of Body Self-Awareness." *Infancy* 11 (2): 157–174.
- Morales, Michael, Peter Mundy, Christine E. F. Delgado, Marygrace Yale, Daniel Messinger, Rebecca Neal, and Heidi K. Schwartz. 2000. "Responding to Joint Attention across the 6- through 24-Month Age Period and Early Language Acquisition." *Journal of Applied Developmental Psychology* 21 (3): 283–298.
- Mundy, Peter, Marian Sigman, and Connie Kasari. 1990. "A Longitudinal Study of Joint Attention and Language Development in Autistic Children." *Journal of Autism and Developmental Disorders* 20 (1): 115–128.
- Murata, Shingo, Jun Namikawa, Hiroaki Arie, Shigeki Sugano, and Jun Tani. 2013. "Learning to Reproduce Fluctuating Time Series by Inferring Their Time-Dependent Stochastic Properties: Application in Robot Learning via Tutoring." *IEEE Transactions on Autonomous Mental Development* 5 (4): 298–310.
- Nagai, Yukie. 2005a. "Learning to Comprehend Deictic Gestures in Robots and Human Infants." In *IEEE International Workshop on Robot and Human Interactive Communication*, 217–222. New York: IEEE.
- Nagai, Yukie. 2005b. "The Role of Motion Information in Learning Human-Robot Joint Attention." In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2069–2074. New York: IEEE.
- Nagai, Yukie. 2019. "Predictive Learning: Its Key Role in Early Cognitive Development." *Philosophical Transactions of the Royal Society B* 374 (1771): 20180030.
- Nagai, Yukie, Minoru Asada, and Koh Hosoda. 2006. "Learning for Joint Attention Helped by Functional Development." *Advanced Robotics* 20 (10): 1165–1181.
- Nagai, Yukie, Koh Hosoda, Akio Morita, and Minoru Asada. 2003. "A Constructive Model for the Development of Joint Attention." *Connection Science* 15 (4): 211–229.
- Nagai, Yukie, Yuji Kawai, and Minoru Asada. 2011. "Emergence of Mirror Neuron System: Immature Vision Leads to Self-Other Correspondence." In Vol. 2, *2011 IEEE International Conference on Development and Learning*, 1–6. New York: IEEE.
- Nagai, Yukie, and Katharina J. Rohlfing. 2009. "Computational Analysis of Motionese toward Scaffolding Robot Action Learning." *IEEE Transactions on Autonomous Mental Development* 1 (1): 44–54.
- Nakajo, Ryoichi, Maasa Takahashi, Shingo Murata, Hiroaki Arie, and Tetsuya Ogata. 2016. "Self and Non-self Discrimination Mechanism Based on Predictive Learning with Estimation of Uncertainty." In *International Conference on Neural Information Processing*, 228–235. Cham, Switzerland: Springer.
- Paulus, Markus. 2014. "The Emergence of Prosocial Behavior: Why Do Infants and Toddlers Help, Comfort, and Share?" *Child Development Perspectives* 8 (2): 77–81.
- Pellicano, Elizabeth, and David Burr. 2012. "When the World Becomes 'Too Real': A Bayesian Explanation of Autistic Perception." *Trends in Cognitive Sciences* 16 (10): 504–510.

- Philippesen, Anja, and Yukie Nagai. 2018. "Understanding the Cognitive Mechanisms Underlying Autistic Behavior: A Recurrent Neural Network Study." *Proceedings of the 8th IEEE International Conference on Development and Learning and on Epigenetic Robotics*, 84–90. New York: IEEE.
- Philippesen, Anja, and Yukie Nagai. 2019. "A Predictive Coding Model of Representational Drawing in Human Children and Chimpanzees." In *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics*, 171–176. New York: IEEE.
- Premack, David, and Guy Woodruff. 1978. "Does the Chimpanzee Have a Theory of Mind?" *Behavioral and Brain Sciences* 1 (4): 515–526.
- Rizzolatti, Giacomo, Leonardo Fogassi, and Vittorio Gallese. 2001. "Neurophysiological Mechanisms Underlying the Understanding and Imitation of Action." *Nature Reviews Neuroscience* 2 (9): 661–670.
- Rizzolatti, Giacomo, Luciano Fadiga, Vittorio Gallese, and Leonardo Fogassi. 1996. "Premotor Cortex and the Recognition of Motor Actions." *Cognitive Brain Research* 3 (2): 131–141.
- Rochat, Philippe. 2003. "Five Levels of Self-Awareness as They Unfold Early in Life." *Consciousness and Cognition* 12 (4): 717–731.
- Scaife, Michael, and Jerome S. Bruner. 1975. "The Capacity for Joint Visual Attention in the Infant." *Nature* 253 (5489): 265–266.
- Schillaci, Guido, Verena V. Hafner, and Bruno Lara. 2016. "Exploration Behaviors, Body Representations, and Simulation Processes for the Development of Cognition in Artificial Agents." *Frontiers in Robotics and AI* 3:39.
- Senju, Atsushi, and Gergely Csibra. 2008. "Gaze Following in Human Infants Depends on Communicative Signals." *Current Biology* 18 (9): 668–671.
- Sommerville, Jessica A., Amanda L. Woodward, and Amy Needham. 2005. "Action Experience Alters 3-Month-Old Infants' Perception of Others' Actions." *Cognition* 96 (1): B1–B11.
- Sumioka, Hidenobu, Koh Hosoda, Yuichiro Yoshikawa, and Minoru Asada. 2007. "Acquisition of Joint Attention through Natural Interaction Utilizing Motion Cues." *Advanced Robotics* 21 (9): 983–999.
- Sumioka, Hidenobu, Yuichiro Yoshikawa, and Minoru Asada. 2008. "Learning of Joint Attention from Detecting Causality Based on Transfer Entropy." *Journal of Robotics and Mechatronics* 20 (3): 378.
- Sumioka, Hidenobu, Yuichiro Yoshikawa, and Minoru Asada. 2010. "Reproducing Interaction Contingency toward Open-Ended Development of Social Actions: Case Study on Joint Attention." *IEEE Transactions on Autonomous Mental Development* 2 (1): 40–50.
- Tani, Jun. 1998. "An Interpretation of the 'Self' from the Dynamical Systems Perspective: A Constructivist Approach." *Journal of Consciousness Studies* 5 (5–6): 516–542.
- Tomasello, Michael, and Michael Jeffrey Farrar. 1986. "Joint Attention and Early Language." *Child Development* 57 (6): 1454–1463.
- Triesch, Jochen, Hector Jasso, and Gedeon O. Deák. 2007. "Emergence of Mirror Neurons in a Model of Gaze Following." *Adaptive Behavior* 15 (2): 149–165.
- Triesch, Jochen, Christof Teuscher, Gedeon O. Deák, and Eric Carlson. 2006. "Gaze Following: Why (Not) Learn It?" *Developmental Science* 9 (2): 125–147.
- Uddin, Lucina Q., Mari S. Davies, Ashley A. Scott, Eran Zaidel, Susan Y. Bookheimer, Marco Iacoboni, and Mirella Dapretto. 2008. "Neural Basis of Self and Other Representation in Autism: An FMRI Study of Self-Face Recognition." *PLoS One* 3 (10): e3526.
- Van de Cruys, Sander, Kris Evers, Ruth Van der Hallen, Lien Van Eylen, Bart Boets, Lee de-Wit, and Johan Wagemans. 2014. "Precise Minds in Uncertain Worlds: Predictive Coding in Autism." *Psychological Review* 121 (4): 649.
- Warneken, Felix, and Michael Tomasello. 2006. "Altruistic Helping in Human Infants and Young Chimpanzees." *Science* 311 (5765): 1301–1303.
- Warneken, Felix, and Michael Tomasello. 2007. "Helping and Cooperation at 14 Months of Age." *Infancy* 11 (3): 271–294.
- Weintraub, Karen. 2011. "Autism Counts." *Nature* 479 (7371): 22.
- Woodward, Amanda L. 1998. "Infants Selectively Encode the Goal Object of an Actor's Reach." *Cognition* 69 (1): 1–34.
- Yamada, Yasunori, Hoshinori Kanazawa, Sho Iwasaki, Yuki Tsukahara, Osuke Iwata, Shigehito Yamada, and Yasuo Kuniyoshi. 2016. "An Embodied Brain Model of the Human Foetus." *Scientific Reports* 6:27893.
- Yoshikawa, Yuichiro, Yoshiki Tsuji, Koh Hosoda, and Minoru Asada. 2004. "Is It My Body? Body Extraction from Uninterpreted Sensory Data Based on the Invariance of Multiple Sensory Attributes." In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Vol. 3. Cat. No. 04CH37566, 2325–2330. New York: IEEE.

19 Human-Robot Interaction

Tony Belpaeme

19.1 Introduction

Human-robot interaction (HRI) studies the interaction between people and robotic systems. While robots are traditionally operated using user interfaces gleaned from human-computer interaction, such as control panels or screen-based interfaces, there is potential to move toward more natural modes of interaction. These will, to a large extent, be modeled on how people interact with each other and are composed of verbal and nonverbal ways of interacting.

HRI is a broad church: at one end of the spectrum, it studies how an operator can control one or more robotic systems through traditional methods and sometimes focuses on the cognitive load imposed by controlling one or more robots. For example, if an operator coordinates a handful of semiautonomous drones during a search and rescue operation, how can the cognitive load on the operator be optimized to maximize the efficiency of the overall mission (e.g., Goodrich et al. 2011)? On the other end of the spectrum of HRI, one finds research into natural interaction between humans and robots. This field is also known as *social robotics*, and the large majority of research efforts in HRI concentrate on it (Bartneck et al. 2020). The holy grail of social HRI, of course, is the natural and intuitive interaction between people and artificial systems. On one hand, this is a technical effort, with results in social signal processing, artificial intelligence, and robotics coming together to create social robots. But social robotics offers a unique opportunity to study how people respond and interact with artificial social agents. Social robots take up a singular position in agents we interact with. The interaction between people has, of course, been the subject of extensive study for more than a century, and the interaction between animals and people has been researched at length, but robots are a new and, until recently, unexplored “species.” Until recently, we have known very little about how people interact with robots, and our relation and interaction with robots is continuously evolving. Culture, media, education, context, and exposure change our attitudes toward robots and the ways in which we interact with them. When we meet a robot, several automatic social responses kick in that color our interaction with the robot; these responses evolved or developed to interact with other humans and often transfer to our interaction with robots.

This is not unique to robots. We treat all technology to some extent as if it is humanlike, something known as anthropomorphization, which Clifford Nass called the “media equation.” We relate to media—computers, printers, mobile phones, and of course robots—as if they are human (Reeves and Nass 1996). Everyone has at one time or another muttered at their computer when it crashed or cursed their printer when the paper jammed, but the media equation theory takes things a little further by claiming that we not only respond to these media as if they were persons but ascribe personal qualities to each, such as a personality, expertise, and even gender. And we often do so without being aware of it. The media equation is taken to the extreme in social robots, as the appearance of the robot and its behavior (the things it does) have been carefully designed to elicit a strong social response from us.

19.2 Cognitive and Neuroscientific Insights Informing HRI

Social psychology is immediately relevant to the design of social robots, and knowingly or not, designers and programmers of social robots take concepts and theories from social psychology into consideration when building robots. Failing to do so usually results in a disappointing HRI. Whether you wish to create a friendly robot or a horror experience, you will rely on fundamentals from social psychology when designing the appearance of your robot and its interaction.

The media equation predicts that people will perceive and treat robots in a humanlike way, but the fact that we readily interpret animated objects as having humanlike emotions and intentions has been known for a long time. Fritz Heider and Marianne Simmel (1944), two psychologists working together in the United States, published an influential paper titled “An Experimental Study of Apparent Behavior” in which they described a simple and elegant experiment: They asked people to describe short film clips of moving geometric figures, such as circles and triangles. The figures were animated by hand and seemed to play out a short story. Everyone who saw the videos ascribed emotions and intentions to the figures. The original videos from the 1940s can still be found online, and even now when seeing the videos, people readily see the figures having emotions, intentions, and motivations, and they see a narrative unfold over the few minutes of video runtime. This is our social brain interpreting the world around it and, specifically, our theory of mind—our ability to attribute mental states to others and ourselves—overinterpreting moving geometric figures. This concept has been gratefully used by animators, and some striking examples exist of very minimalist animation films that show that very little is needed to nudge our social brain into interpreting simple shapes and movement as having agency (Thomas and Johnston 1995). If you have ever observed a vacuuming robot moving around the room, you have probably been struck by its animallike appearance as it scuttles around the room, gently bumping into furniture and working hard at getting specks of dirt from the floor. These robots are not designed to be social, and yet they still evoke a strong social response in us. In social robots, designers add elements such as a head, eyes, and reactive responses to evoke a strong social response in people.

One such social response on which designers rely is *pareidolia*: the tendency to see human or animal forms in objects, such as dogs in clouds or the face of Elvis on a piece

of burnt toast. Using magnetoencephalography (MEG), researchers found that the ventral fusiform face area (FFA) in the brain is involved. The FFA has been implicated in detecting faces of people and animals and is also involved in distinguishing animate from inanimate visual stimuli (Kanwisher et al. 1999). This area shows a cortical response 170 ms after we are presented with a human face and shows a similar but slightly earlier activation of 165 ms when seeing objects that resemble faces (Hadjikhani et al. 2009). This suggests that seeing faces is a very early and automatic response and is not something the brain puzzles together after extended cognitive processing. As such, we can assume that responses to robots with a face are early and automatic.

19.3 Design of Social Robots

One aspect that often arises in robot design is that of neoteny, a juvenile appearance that usually evokes a caring response and is generally described as “cute.” Young animals, including human children, have a large head, large eyes, chubby cheeks, a small chin, a flat face, a small nose, and relatively short arms and legs. Konrad Lorenz (1982) argued that infantile and juvenile features have a biological function by triggering nurturing responses in adults. We are so keen on neotenous appearances that we breed domesticated animals to retain neotenous features. Many breeds of smaller dogs retain juvenile features, such as a short snout and a relatively large head and large eyes, and consequently are considered cute by most people. The nurturing response is also largely cross-cultural. The same physical features evoke a similar response in people regardless of culture or background. This has been used to good effect by robot designers: if a robot is to be likeable, designers will give it features that evoke a caring response. This not only causes people interacting with the robot to find it cute but also makes them inclined to feel more generous toward any mistakes the robot makes. The opposite seems to hold as well. Robots that have adult, or gerontomorphic, features appear less cute and have less appeal. While there is no research on this yet, it is likely that they are considered more knowledgeable and authoritative, and therefore it makes sense for robot designers to give robots that need to radiate authority or trust an adult appearance (see figure 19.1).

Perhaps the most well-known issue in robot design is that of the *uncanny valley* (figure 19.2). This effect, first hypothesized by Mori in 1970 (Mori et al. 2012), describes the familiarity or appeal of a robot as a function of its human likeness. Mori in his original paper wrote about 親和感 (*shinwa-kan*), which does not translate well into English but is sometimes described as familiarity, appeal, likeability, or affinity. When a robot does not resemble a human, it has low familiarity. This gradually goes up: As human likeness increases, so does familiarity, until the robot is almost humanlike but not quite. At this point familiarity gets knocked back, and when plotted this resembles a sharp dip in the familiarity curve. This is known as the uncanny valley. Androids, robots that have humanlike skin but lack humanlike motions, find themselves firmly in the uncanny valley. You can climb out of the uncanny valley by making a robot that is almost indistinguishable from a person. Note that the uncanny valley effect is more pronounced when the robot is moving: the familiarity or eeriness of the robot is more exaggerated when the robot is animated. Mori never backed up his hypothesis with data, but later empirical research has shown that the




Figure 19.1

A neotenus appearance, characterized by a large forehead, big eyes, a small mouth, and a large head, in robots such as the SoftBank Robotics NAO robot (*left*), make people feel more attracted to them. Robots with adultlike features, such as the Engineered Arts SociBot, which has an adult face (*right*), are likely to be found more authoritative and knowledgeable.

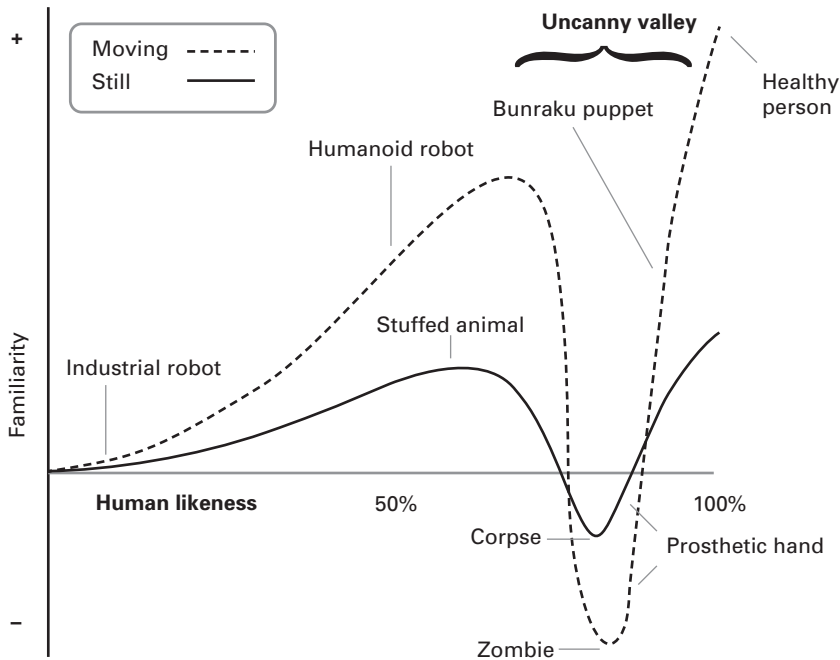


Figure 19.2

A plot showing the uncanny valley, with the famous dip when robots look almost humanlike but repel us because they are not sufficiently humanlike. *Source:* Based on Mori 1970, Wikimedia.

uncanny valley is indeed real (MacDorman and Ishiguro 2006; MacDorman and Chattopadhyay 2016).

Rosenthal-von der Pütten et al. (2019) studied the neural mechanisms underlying human responses to artificial agents and, specifically, the uncanny valley response. They suggest that the uncanny valley requires a neural system that derives human likeness from sensory cues followed by a downstream system that integrates these signals into a nonlinear value function representing the uncanny valley response curve. Using functional magnetic resonance imaging (fMRI), they investigated the neural activity of people when observing people and artificial agents, including robots, while making rated responses or expressing a preference for stimuli. They found that the ventromedial prefrontal cortex encoded a representation of the uncanny valley, in which the subjective likability of artificial agents was a nonlinear function of human likeness. Functionally connected areas in the brain encoded critical inputs for signals: the temporoparietal junction (TPJ) encoded a linear human likeness continuum. The TPJ was also found to be active in detecting agency (Mar et al. 2007), belief attribution, and learning from others (Rosenthal-von der Pütten et al. 2019). In addition, nonlinear representations of human likeness found in the dorsomedial prefrontal cortex (DMPFC) and fusiform gyrus (FFG) emphasized a human-nonhuman distinction. The DMPFC is known to show activity when attributing mental states to others or when assessing performance of others or of the self (Rosenthal-von der Pütten et al. 2019), while the FFG is implicated in distinguishing animate from inanimate stimuli (Chaminade et al. 2010). Activation in the amygdala, which in humans is implicated in the formation and storage of memories associated with emotional events, was found to predict a negative response to artificial agents. As such, the brain seems to have a direct neural representation of the uncanny valley, or rather the uncanny valley can be explained by brain processes that are universal to all people.

If the same neural mechanisms implicated in assessing people, people's behavior, and the agency of stimuli are also active when we perceive robots, then this might help us design more effective robots. Generally, what makes people appealing will make robots appealing, and only cultural conditioning and habituation are likely to change the initial, and often automatic, responses we have to robots.

When discussing the uncanny valley, one cannot escape mentioning androids and perhaps their more famous ilk, the Geminoids. A Geminoid—a contraction of Gemini (meaning “twins” in Latin) and android—is modeled after a human being and as such is their robotic doppelgänger. Hiroshi Ishiguro was the first to build Geminoids, and the various models that have been built—including ones of himself, his daughter, and a Japanese news anchor—have been the subject of academic study into the uncanny valley effect. These studies showed that the uncanny valley effect is sometimes not there or cannot be explained by relying on appearance alone. Bartneck et al. (2009) had people briefly interact with Hiroshi Ishiguro or with his Geminoid. While participants could clearly distinguish an android from a human, and unsurprisingly found the human to be more humanlike, the android was not liked less, which goes against Mori's prediction. This result and others suggest that the uncanny valley is a multidimensional phenomenon and that the two-dimensional plot of figure 19.2 should be revised. Instead the effect is caused by a mismatch between different aspects of the robot: a robot that appears human but moves like a robot causes tension in the observer, which leads to an eerie appearance (Moore 2012).




Figure 19.3 Hiroshi Ishiguro and his Geminoid, a robot replica used to study people’s responses to lifelike robots. *Source:* Osaka University, Intelligent Robotics Laboratory.

19.4 Verbal Interaction

Social robots will often be addressed using language. Even robots that are not humanlike in appearance, such as animallike robots, are often addressed using speech. Depending on the robot’s appearance, people might expect a coherent linguistic response. We don’t expect a robot dinosaur to talk back, but we do have expectations of humanoid robots and are invariably somewhat disappointed when those expectations are not met.

In addition, language is most likely to be the most natural and therefore intuitive way to interact with robots. But despite the use of language seeming effortless to us, verbal interaction between people and robots is still a formidable challenge. The typical approach in building natural language interaction (NLI) has been to cut up the problem into several components: speech recognition, dialogue management, language generation, and speech production. And while progress is being made in each of these, unconstrained natural language interaction is still well beyond our technical grasp. Speech recognition, using deep neural networks trained on large sets of annotated speech, now performs better than human transcribers for English spoken by adults (e.g., Xiong et al. 2018). Speech production is almost indistinguishable from human speech for the reading of text with neutral prosody (van den Oord et al. 2016). The developments in speech recognition and speech production have led to a raft of novel applications. Prime examples are the digital assistants, such as Amazon’s Alexa or Apple’s Siri assistants, that can act on spoken instructions and respond using speech. But these assistants are still very much limited in their func-

tionality, as are most spoken NLI applications. They can take short phrases and take the user through a turn-based dialogue to fill in slots, but they cannot engage in unconstrained dialogue. They do struggle with pragmatic language use—that is, the social language that we use in our daily interactions with others, from the short utterances such as “yup,” “sure,” or “dunno” that keep linguistic interaction flowing to the extensive reliance on contextual cues to interpret and produce linguistic utterances.

When comparing artificial linguistic interaction systems to language processing in the human brain, it is clear that the two are far apart on several levels. At a fundamental level, language in computers is meaningless to the computer. A chatbot can utter phrases about feelings or the weather, but it does not really understand what it is talking about. It has never experienced feelings or weather, or any other words for that matter. The words that a chatbot uses are not *grounded*. Grounding happens when words and linguistic expressions are experienced and from that become meaningful. The word “chair” only becomes meaningful when a computer or robot has an experiential sensation of a chair by seeing a chair through its camera, by feeling a chair through tactile sensors, or by understanding the function of a chair.

There have been some interesting developments in statistical language processing, where algorithms are used to build models of a language by analyzing large corpora of text. The earliest such algorithms built cooccurrence statistics of words, basically counting which words appeared near others in texts. A distance measure is used to report which words are closer in meaning and which are not. One such technique, latent semantic analysis (LSA), can tell that “king” and “queen” are closely related and that “king” and “lemon” are not (Landauer et al. 1998). New neural network-based approaches take statistical cooccurrence further by learning long-distance dependencies between words. The most recent solutions use recurrent neural networks. At the time of writing, the most notable model is the generative pretrained transformer 3, or GPT-3, but given the arms race between large corporations to outperform each other’s language models, the GPT-3 will soon be superseded. The GPT-3 uses transformer networks and was trained on hundreds of billions of words. It was tasked with learning to predict the next word in a sentence and by doing so built a model not only of the English language but also of programming languages (Brown et al. 2020). The GPT-3 seems to have a firm grasp on semantics. It can not only complete sentences; there are impressive examples of it completing short-story lines starting from only an opening paragraph. It can answer questions and passes tests aimed at assessing the vocabulary skills of children. From a cursory inspection, it would seem that the GPT-3 understands language, as it uses language in a very coherent way. However, while the GPT-3 can tell you who the president of the United States is, it would not be able to recognize the president in a photo. The reason, of course, is that the GPT-3, and all other text-based natural language processing systems, are completely text based: the words they use are not *grounded*.

The contrast with human cognition could not be greater: all the words and linguistic constructions we use are grounded in a sensory reality (Harnad 1990). Many have argued that robots should do the same if they are to interact with people in a way in which our exchanges are meaningful (Cangelosi et al. 2002). A robot without grounded linguistic symbols can seem to know the “color of grass,” but if it is not able to tie the visual perception of green and grass together, together with all the other memories and cultural agreements on language, human-robot conversation is likely to remain fairly limited.

Another challenge, especially in the context of cognitive robotics, is that language in the human brain is rather poorly understood. We can prod the linguistic brain through behaviorist experiments—for example, by measuring response times to words, which gives us an insight into how words and their meaning might be represented in the brain. Or we sometimes get intriguing views into the linguistic brain through patients who have suffered brain injuries. Important brain regions implicated in language processing and production, such as Broca's and Wernicke's areas, were discovered after studying patients with lesions to those areas. We also discovered that language is to some extent processed in the right hemisphere, after studying patients who had both hemispheres separated by cutting the *corpus callosum*, the part of the brain connecting both hemispheres, but were still able to interpret words shown to only the right visual field.

But even modern brain-imaging techniques have shed relatively little light on how language is processed (Dronkers et al. 2004), represented (Hagoort 2005), and produced in the brain (Levelt 2001) and certainly not to an extent in which insights from cognitive neuroscience would enable us to build better natural language interaction systems. If there is perhaps one valuable lesson, it is that language is not compartmentalized. Instead language seems to permeate the entire brain, with some clear loci for more specific language functions. Artificial NLP, on the other hand, is compartmentalized into components such as speech recognition, language interpretation, dialogue processing, language generation, and speech production while ignoring elements often essential to linguistic communication. Most importantly, the multimodal and nonverbal aspects of communication are largely ignored, and artificial NLP is therefore rather impoverished. Two examples should make this clear: prosody and priming. Prosody is ignored in NLP, although the meaning of a spoken utterance can be completely changed through prosody. Just think of the many ways in which “I'm not at all angry” can be expressed and how the meaning of such a short sentence can swing between joking, furious, irritated, or sad. Human linguistic perception and production is fine-tuned for this, but it remains firmly outside the grasp of artificial speech recognition and production.

Priming is the effect whereby one stimulus influences the response to a later stimulus. For example, asking, “What do cows drink?” often results in people answering “milk” instead of “water” (Rose et al. 2015). Language in the brain is organized as an associative network, with sounds, words (or lemmas), and meaning connected in networks (Collins and Loftus 1975; Levelt 2001). Statistical methods of language modeling, such as hidden Markov models or long short-term memory networks, indispensable in speech recognition and machine translation, explicitly learn statistical associations between phonemes and words. Priming is a very important mechanism both in the brain and in these artificial models: the presentation of a word or phoneme primes, or rather predicts, the next most probable word or phoneme. In the brain, priming is multimodal (Wood et al. 2012), but in NLP the priming only happens within the phonetic or lexical domain, thereby cutting NLP off from modalities that the human brain relies upon to disambiguate and enrich language.

19.5 Nonverbal Interaction

Most content of a natural interaction is contained in its nonverbal aspects. Of course, written text contains very little nonverbal communication (apart from the occasional emoticon) and

seems to work well at conveying information. But spoken language, and specifically language spoken in the presence of others, relies heavily on nonverbal elements. The division of labor between verbal and nonverbal is contested. A widely cited statement is that of Mehrabian (1972), which claims that 55 percent of communication is contained in body language, 38 percent in tone of voice, and only 7 percent in the words spoken. While the exact ratio is up for debate, the fact that verbal communication only accounts for a fraction of communication should point out the flaws in our current efforts in building HRI. For historical reasons most of our technical efforts have been on creating verbal or text-based linguistic interactions while at the same time ignoring nonverbal aspects of interaction. And if we did study nonverbal interaction, we studied it in isolation from other communication channels.

Emotion is a textbook example of this: Due to technical and resource limitations, the first studies of emotion used photographs of facial expressions. Paul Ekman, in his effort to show that some emotions are universal, took a number of photographs of himself and others showing extreme emotions, such as happiness or anger. He indeed confirmed that these emotions are universally recognized and, building on this work, argued that there are at least six or seven basic emotions (Ekman 1972, 1992). Ekman built on a tradition started by Darwin (1872) of using photographs of faces to study emotions, and ever since the discussion of emotions has been dominated by a focus on facial expressions. Nevertheless, faces only show extreme emotions, and emotion is much more likely to be gleaned from context and other body cues (Kappas 2003). In a striking experiment, it was shown that the body posture of tennis players, rather than their facial expressions, showed whether they had won or lost a point, convincingly demonstrating that the face is not necessarily a window to the soul, or to emotions in this case (Aviezer et al. 2012).

Just as with anthropomorphization, the human brain is ever eager to interpret nonverbal signals as meaningful. The clicks, beeps, and whirrs that R2-D2, one of the robot leads from the *Star Wars* series, emits are never interpreted as background noise on the soundtrack of the film but are interpreted as meaningful and relevant by the cinema audience. These clicks and beeps, or nonlinguistic utterances (NLU), can be used to add a nonlinguistic communication channel to robots, complementing language or even short-cutting the need for language. NLUs are interpreted as meaningful by children and adults and can be used to communicate the emotional state of the robot (Read and Belpaeme 2014; see figure 19.4).

Further analysis showed how NLUs are interpreted categorically: if people are asked to interpret an NLU as an emotion, then their interpretation is being drawn to one of only a handful of basic emotions such as happiness, anger, surprise, or fear (Read and Belpaeme 2016). Categorical perception is a fundamental property of perception and is instrumental in interpreting perceptual stimuli. The human brain interprets sensory perception as belonging to a limited number of conceptual states. For example, speech sounds are interpreted as belonging to only a distinct number of phonemes. If hearing a speech continuum in which the amount of voicing is changed gradually, from not at all in “p” in /pa/ to fully voiced in “b” in /ba/, then the perception will be drawn toward known vowels, either “pa” or “ba” but nothing in between. It is surprising that the cognitive mechanisms used to interpret human-human verbal and nonverbal communication are still at work when we are interpreting robotic communicative signals.




Figure 19.4

Random robot sounds, a concatenation of clicks and beeps, were played to children between six and eight years old. The children were asked to show which emotion the robot was displaying by recreating the emotion on a digital face. These responses were then mapped to a 3D emotion space. Instead of responses being uniformly scattered over the emotion space plot, the children's responses clustered together near basic emotion. This suggests that robot sounds are interpreted as humanlike emotions and that this process is categorical. *Source:* From Read and Belpaeme 2012, 2016.

The combination of verbal and nonverbal interaction, often referred to as multimodal communication in technical parlance, is perhaps the biggest challenge in HRI. One of the reasons for this is that a divide-and-conquer approach, in which a problem is divided up into smaller problems, each to be solved on their own before being recombined to form a total solution, does not seem promising when it comes to building multimodal HRI. In human cognition, multimodal interaction is a complex activity to which all cognitive faculties contribute without clear division, sequence, or hierarchy. For example, hearing a verb (such as “kick”) activates the corresponding action in the motor cortex (activity when kicking or thinking about kicking; Pulvermüller 1999), and hearing a naturalistic sound (such as a dog's “woof”) and spoken words (/dɔg/) 346 ms before a picture search task led to faster visual detection of the picture of a dog from between distractors (Chen and Spence 2011). It is very likely that the cognitive organization of human interaction will need to be reflected in some way when building HRI. The current separation of processing, with separate components such as speech recognition, dialogue, text to speech, emotion recognition, facial expressions, gesture production, or prosody is artificial and does not have the tight and dynamic coupling that is likely to be necessary for natural HRI.

19.6 Applications

A better understanding of the cognitive mechanisms involved in HRI would surely allow us to build better robots, better interactions, and the best applications. For now, the design of robots and interactions has relied a lot on the gut feeling of designers and engineers and to a lesser extent on theory. However, as soon as HRI is used for applications, an improved understanding of the responses of the human brain to robots might be essential.

Social robots can be used to entertain, persuade, and inform. The strong social character of robots lends itself well to establishing a social bond, and this can be used in diverse applications, such as retail, education, or therapy.

Robots show potential in education. When compared to screen-based learning technologies, such as educational software on computers or tablets, robots tend to have better outcomes. This can be explained by the explicit and tangible social character of the robots, which leads to both improved attitudes toward learning and better learning outcomes. In a metareview (Belpaeme et al. 2018), papers comparing tutoring robots against an alternative, such as educational software or an on-screen avatar, showed that the mean cognitive outcome effect size (Cohen's d) of robot tutoring is 0.70 (95 percent confidence interval (CI), 0.66 to 0.75), which compares favorably to what human tutors can achieve: human tutors achieve an outcome effect size of $d=0.79$ (Vanlehn 2011). While robot tutors do show promise, designing a robot tutor still is challenging. Robots can be used to tutor restricted domains, such as simple math exercises, but little is known about how to design robot tutors that tackle harder learning challenges. One such challenge is language: the current school-based teaching of a second language relies a great deal on class-based learning of vocabulary and grammar with little to no attention to language use and interaction. This is far removed from how a first language is seemingly effortlessly acquired through interacting with parents, siblings, and peers. The main reason why school-based language learning is so different is that the teacher cannot engage in interaction on an individual basis with all pupils in the classroom. And this is where robots show considerable promise: a robot has the time and infinite patience to interact with those learning a target language. A robot probably also has a better accent than the teacher and can personalize its tutoring to the learner.

Vogt et al. (2019) reported on a large-scale study in which a language-tutoring robot helped young children learn the words and grammar of a second language (see figure 19.5). They used a NAO robot to teach English to five-to-six-year-olds in the Netherlands. Children learned not only nouns (“giraffe” or “boy”) but also words used in numeracy (counting words or quantities, such as “more” or “fewer”) and spatial language (such as “behind,” “in front,” and “next to”). The robot tutored the children over seven lessons, introducing six new words during every lesson. The study was used not only to establish whether the robot would be better than only a tablet but also to see whether a robot using gestures to accentuate the words would be a better language tutor. It was divided over four study conditions (a control condition receiving no tutoring, a tablet-only condition, a robot without gestures condition, and a robot with gestures condition), and 208 children took part. While the children did learn English, no significant difference could be found between the learning outcomes: children did not learn more from a robot, whether it was using gestures or not, than from a tablet alone. While there are demonstrations of robots being very effective tutors in narrow domains, the



Figure 19.5
A child learning a second language with the support of a social robot.

benefits of using robots in more complex domains, such as second-language tutoring, are harder won. Robots have been shown to be effective in tutoring vocabulary (van den Berghe et al. 2019), but a more complex use of language probably requires a more complex HRI. A better understanding of how children and adults learn, and how robots can have an impact on this process, will be necessary. It is likely that the social and physical presence of robots is a strong influence on the learning process, but without more open-ended natural interaction, the use of robot tutors is likely to be limited to narrow and closed domains, such as math exercises or vocabulary.

Another application of HRI in which robots are likely to have a significant impact in the future is therapy (Belpaeme et al. 2013). In the last two decades, robotics has been promoted as a promising new technology in autism spectrum disorder (ASD) therapy (Scassellati, Admoni, and Mataric 2012; Thill et al. 2012), and while many supportive case studies exist, there has been a dearth of quantitative empirical evidence about the efficacy of robot therapy (Diehl et al. 2012; Pennisi et al. 2016) that only recently is being resolved. The effect of robots and their behavior on people with ASD is only being studied through the lens of psychological therapy, with little consideration for the cognitive processes involved in the perception of and interaction with robots. It is very likely that a better understanding of the neuropsychology and cognition involved in HRI will allow us to build more effective HRI.

19.7 Conclusion

The relation between human cognition and HRI has largely been explored at the behavioral level. Recently, brain-imaging techniques and response time experiments have given us a

view on how the brain responds to robot stimuli and interactions with robots. All data seem to suggest that interaction with robots relies on the very same social cognitive mechanisms and neural pathways that are also active when we interact with people. This in itself is not very surprising: the brain just generalizes, and our social cognition spills over to nonhuman agents, be they pets or robots. What is more surprising is that our brain readily interprets robotic behaviors, robot forms, and robot noises for which our brain certainly did not evolve. Of course, the nonlinguistic utterances of fictional robots and toy robots have been designed to be interpretable, but even odd combinations—such as a robot vacuum cleaner with a wagging tail (Singh and Young 2012)—remain legible and socially meaningful to us, showing that the human brain really is a most gregarious social interpreter. Understanding how it accomplishes that is likely to lead to a more efficient design of new forms and behavior in HRI.

Additional Reading and Resources

- A classic survey of early approaches to HRI: Goodrich, Michael A., and Alan C. Schultz. 2007. “Human-Robot interaction: A Survey.” *Foundations and Trends in Human-Computer Interaction* 1 (3): 203–275.
- A recent, comprehensive volume on HRI: Bartneck, Christoph, Tony Belpaeme, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic. 2020. *Human-Robot Interaction: An Introduction*. Cambridge: Cambridge University Press.
- A recent collection on research methods in HRI: Jost, Céline, Brigitte Le Pévédic, Tony Belpaeme, Cindy Bethel, Dimitrios Chrysostomou, Nigel Crook, Marine Grandgeorge, and Nicole Mirnig, eds. 2020. *Human-Robot Interaction: Evaluation Methods and Their Standardization*. Vol. 12. Berlin: Springer.
- A time line of HRI, podcasts on HRI, and additional material accompanying Bartneck et al. (2020): <https://www.human-robot-interaction.org/>.
- The portal link to the flagship HRI conference in the field and resources on HRI: <http://humanrobotinteraction.org/>.
- A one-hour video introduction to HRI and social robotics: <https://www.youtube.com/watch?v=Lpp1FjkOyN4>.

References

- Aviezer, Hillel, Yaacov Trope, and Alexander Todorov. 2012. “Body Cues, Not Facial Expressions, Discriminate between Intense Positive and Negative Emotions.” *Science* 338 (6111): 1225–1229.
- Bartneck, Christoph, Tony Belpaeme, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic. 2020. *Human-Robot Interaction: An Introduction*. Cambridge: Cambridge University Press.
- Bartneck, Christoph, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2009. “My Robotic Doppelgänger—a Critical Look at the Uncanny Valley.” In *The 18th IEEE International Symposium on Robot and Human Interactive Communication*, 269–276. New York: IEEE.
- Belpaeme, Tony, Paul Baxter, Robin Read, Rachel Wood, Heriberto Cuayahuitl, Bernd Kiefer, Stefania Racioppa, et al. 2013. “Multimodal Child-Robot Interaction: Building Social Bonds.” *Journal of Human-Robot Interaction* 1 (2): 33–53.
- Belpaeme, Tony, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. 2018. “Social Robots for Education: A Review.” *Science Robotics* 3 (21).

- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models are Few-Shot Learners." ArXiv preprint: 2005.14165.
- Cangelosi, Angelo, A. Greco, and S. Harnad. 2002. "Symbol Grounding and the Symbolic Theft Hypothesis." In *Simulating the Evolution of Language*, edited by A. Cangelosi and D. Parisi, 191–210. London: Springer.
- Chaminade, Thierry, Massimiliano Zecca, Sarah-Jayne Blakemore, Atsuo Takanishi, Chris D. Frith, Silvestro Micera, Paolo Dario, Giacomo Rizzolatti, Vittorio Gallese, and Maria Alessandra Umiltà. 2010. "Brain Response to a Humanoid Robot in Areas Implicated in the Perception of Human Emotional Gestures." *PLoS One* 5 (7): e11577.
- Chen, Yi-Chuan, and Charles Spence. 2011. "Crossmodal Semantic Priming by Naturalistic Sounds and Spoken Words Enhances Visual Sensitivity." *Journal of Experimental Psychology: Human Perception and Performance* 37 (5): 1554.
- Collins, Allan M., and Elizabeth F. Loftus. 1975. "A Spreading-Activation Theory of Semantic Processing." *Psychological Review* 82 (6): 407.
- Darwin, C. 1872. *The Expression of the Emotions in Man and Animals*. London: John Murray.
- Dennett, Daniel C. 1996. *The Intentional Stance*. 6th ed. Cambridge, MA: MIT Press.
- Diehl, J. J., Schmitt, L. M., Villano, M., and Crowell, C. R. 2012. "The Clinical Use of Robots for Individuals with Autism Spectrum Disorders: A Critical Review." *Research in Autism Spectrum Disorders* 6 (1): 249–262.
- Dronkers, Nina F., David P. Wilkins, Robert D. Van Valin Jr., Brenda B. Redfern, and Jeri J. Jaeger. 2004. "Lesion Analysis of the Brain Areas Involved in Language Comprehension." *Cognition* 92 (1–2): 145–177.
- Ekman, Paul. 1972. "Universals and Cultural Differences in Facial Expressions of Emotions." In *Nebraska Symposium on Motivation*, edited by J. Cole, 207–282. Lincoln: University of Nebraska Press.
- Ekman, Paul. 1992. "An Argument for Basic Emotions." *Cognition and Emotion* 6 (3–4): 169–200.
- Goodrich, Michael A., Brian Pendleton, P. B. Sujit, and José Pinto. 2011. "Toward Human Interaction with Bio-inspired Robot Teams." In *2011 IEEE International Conference on Systems, Man, and Cybernetics*, 2859–2864. New York: IEEE.
- Hadjikhani, Nouchine, Kestutis Kveraga, Paulami Naik, and Seppo P. Ahlfors. 2009. "Early (N170) Activation of Face-Specific Cortex by Face-Like Objects." *Neuroreport* 20 (4): 403.
- Hagoort, Peter. 2005. "On Broca, Brain, and Binding: A New Framework." *Trends in Cognitive Sciences* 9 (9): 416–423.
- Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D: Nonlinear Phenomena* 42 (1–3): 335–346.
- Heider, Fritz, and Marianne Simmel. 1944. "An Experimental Study of Apparent Behavior." *American Journal of Psychology* 57 (2): 243–259.
- Kanwisher, Nancy, Damian Stanley, and Alison Harris. 1999. "The Fusiform Face Area Is Selective for Faces Not Animals." *Neuroreport* 10 (1): 183–187.
- Kappas, Arvid. 2003. "What Facial Activity Can and Cannot Tell Us about Emotions." In *The Human Face*, 215–234. Boston: Springer.
- Landauer, Thomas K., Peter W. Foltz, and Darrell Laham. 1998. "An Introduction to Latent Semantic Analysis." *Discourse Processes* 25 (2–3): 259–284.
- Levelt, Willem J. M. 2001. "Spoken Word Production: A Theory of Lexical Access." *Proceedings of the National Academy of Sciences* 98 (23): 13464–13471.
- Lorenz, Konrad. 1982. *The Foundations of Ethology: The Principal Ideas and Discoveries in Animal Behavior*. New York: Simon and Schuster.
- MacDorman, Karl F., and Debaleena Chattopadhyay. 2016. "Reducing Consistency in Human Realism Increases the Uncanny Valley Effect; Increasing Category Uncertainty Does Not." *Cognition* 146:190–205.
- MacDorman, Karl F., and Hiroshi Ishiguro. 2006. "The Uncanny Advantage of Using Androids in Cognitive and Social Science Research." *Interaction Studies* 7 (3): 297–337.
- Mar, Raymond A., William M. Kelley, Todd F. Heatherton, and C. Neil Macrae. 2007. "Detecting Agency from the Biological Motion of Veridical vs Animated Agents." *Social Cognitive and Affective Neuroscience* 2 (3): 199–205.
- Marcus, Aaron, Masaaki Kurosu, Xiaojuan Ma, and Ayako Hashizume. 2017. *Cuteness Engineering: Designing Adorable Products and Services*. Berlin: Springer.
- Mehrabian, Albert. 1972. *Nonverbal Communication*. New York: Routledge.
- Moore, Roger K. 2012. "A Bayesian Explanation of the 'Uncanny Valley' Effect and Related Psychological Phenomena." *Scientific Reports* 2 (2): 864.

- Mori, Masahiro, Karl F. MacDorman, and Norri Kageki. 2012. "The Uncanny Valley." *IEEE Robotics and Automation Magazine* 19 (2): 98–100.
- Pennisi, Paola, Alessandro Tonacci, Gennaro Tartarisco, Lucia Billeci, Liliana Ruta, Sebastiano Gangemi, and Giovanni Pioggia. 2016. "Autism and Social Robotics: A Systematic Review." *Autism Research* 9 (2): 165–183.
- Pulvermüller, Friedemann. 1999. "Words in the Brain's Language." *Behavioral and Brain Sciences* 22 (2): 253–279.
- Read, Robin, and Tony Belpaeme. 2012. "How to Use Non-linguistic Utterances to Convey Emotion in Child-Robot Interaction." In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction*, 219–220. New York: IEEE.
- Read, Robin, and Tony Belpaeme. 2014. "Situational Context Directs How People Affectively Interpret Robotic Non-linguistic Utterances." In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction*, 41–48. New York: IEEE.
- Read, Robin, and Tony Belpaeme. 2016. "People Interpret Robotic Non-linguistic Utterances Categorically." *International Journal of Social Robotics* 8 (1): 31–50.
- Reeves, Byron, and Clifford Ivar Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge: Cambridge University Press.
- Rose, Sebastian Benjamin, Katharina Spalek, and Rasha Abdel Rahman. 2015. "Listening to Puns Elicits the Co-activation of Alternative Homophone Meanings during Language Production." *PLoS One* 10 (6): e0130853.
- Rosenthal-von der Pütten, Astrid, Nicole Krämer, Stefan Maderwald, Matthias Brand, and Fabian Grabenhorst. 2019. "Neural Mechanisms for Accepting and Rejecting Artificial Social Partners in the Uncanny Valley." *Journal of Neuroscience* 39 (33): 6555–6570.
- Scassellati, Brian, Henny Admoni, and Maja Matarić. 2012. "Robots for Use in Autism Research." *Annual Review of Biomedical Engineering* 14:275–294.
- Singh, Ashish, and James E. Young. 2012. "Animal-Inspired Human-Robot Interaction: A Robotic Tail for Communicating State." In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction*, 237–238. New York: IEEE.
- Thill, Serge, Cristina A. Pop, Tony Belpaeme, Tom Ziemke, and Bram Vanderborght. 2012. "Robot-Assisted Therapy for Autism Spectrum Disorders with (Partially) Autonomous Control: Challenges and Outlook." *Paladyn* 3 (4): 209–217.
- Thomas, Frank, and Ollie Johnston. 1995. *The Illusion of Life: Disney Animation*. New York: Hyperion.
- van den Berghe, Rianne, Josje Verhagen, Ora Oudgenoeg-Paz, Sanne van der Ven, and Paul Leseman. 2019. "Social Robots for Language Learning: A Review." *Review of Educational Research* 89 (2): 259–295.
- van den Oord, Aaron, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. "Wavenet: A Generative Model for Raw Audio." ArXiv preprint: 1609.03499.
- Vanlehn, Kurt. 2011. "The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems." *Educational Psychologist* 46 (4): 197–221.
- Vogt, Paul, Rianne van den Berghe, Mirjam De Haas, Laura Hoffman, Junko Kanero, Ezgi Mamus, Jean-Marc Montanier, et al. 2019. "Second Language Tutoring Using Social Robots: A Large-Scale Study." In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction*, 497–505. New York: IEEE.
- Wood, Rachel, Paul Baxter, and Tony Belpaeme. 2012. "A Review of Long-Term Memory in Natural and Synthetic Systems." *Adaptive Behavior* 20 (2): 81–103.
- Xiong, Wayne, Lingfeng Wu, Fil Allewa, Jasha Droppo, Xuedong Huang, and Andreas Stolcke. 2018. "The Microsoft 2017 Conversational Speech Recognition System." In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 5934–5938. New York: IEEE.

20 Language and Communication

Angelo Cangelosi and Tetsuya Ogata

20.1 Introduction

Communication is a rich multimodal process combining spoken language with a variety of nonverbal behaviors such as gaze, gestures, tactile interaction, and emotional cues (Mavridis 2015; Cangelosi and Ogata 2019; Liu and Zhang 2019). For cognitive robotics and human-robot interaction, linguistic and nonverbal communication skills are fundamental cognitive capabilities necessary to interact with people. To ask a humanoid robot to perform a specific task, or to engage in a dialogue with a social robot companion, both people and robots must possess a language-like communication system. In cognitive robotics, the design of speech and nonverbal communication skills is directly inspired by communication in people.

The organization of human language and communication has been the focus of attention in linguistics and psychology. Specific levels of representation and analyses of linguistics skills, ranging from the processing of low-level phonetic features to higher-level communicative and pragmatic processes, have been identified to study language. In addition, developmental psychology has significantly contributed to the identification of the developmental stages and language-learning principles. This has been contextualized within the debate of nativist versus constructivist theories—that is, language acquisition theories giving emphasis to a genetic predisposition to language-related competence versus developmental theories stressing the role of environmental factors. These linguistics and psychology analyses have significantly contributed to the design of cognitively inspired language and communication skills in cognitive robots. Below, we first look at the developmental theories of language learning and the linguistics approach of natural language processing (NLP) and the five levels of analysis. This will inform the discussion of the different models of language acquisition in developmental robotics, of NLP models used in robots, and of the more recent machine-learning models.

20.1.1 Language Development and Learning in Humans

An important issue in language development research is the “nature” versus “nurture” debate. This is the debate between the “nativists,” who hypothesize that babies are born with language-specific knowledge and skills, and the “empiricists,” who propose that

babies construct linguistic knowledge through interaction with their social, language-speaking community. Within the nativist position, influential theories have proposed that there are universal syntactic rules and generative grammar principles (e.g., Chomsky's "brain organ" and "language acquisition device" hypothesis) and that these are innate in the human brain (Chomsky 1965). On the contrary, according to the nurture stance, the essence of linguistic knowledge emerges from language use during development, without any need to assume the existence of innate language-specific knowledge. This empiricist view of language development is also known as the constructivist, usage-based theory of language development (Tomasello 2003; MacWhinney 1998). The child is seen as an active constructor of their own language system through the implicit observation and learning of statistical regularities and logical relationships between the meaning of words and the words used (e.g., cognitive linguistic theories of Goldberg [2006]).

In developmental psychology research, the most significant phenomena of language acquisition occur during the first four years. The early milestones of language development follow the parallel and intertwined development of incremental phonetics-processing capabilities, increasing lexical and grammatical repertoires, and refined communicative and pragmatic faculties. Table 20.1 provides an overview of the main milestones of language development (Hoff 2013; Cangelosi and Schlesinger 2015).

In the first year, the most evident sign of linguistic development concerns phonetic capabilities such as vocal babbling. Babbling initially consists of vocal play with sounds such as cooing, squeals, and growls ("marginal babbling") and later consists of the repetition of language-like syllabic sounds such as "dada" or "bababa" (canonical/reduplicated babbling). Toward the end of the first year, children also start to produce communicative gestures (e.g., pointing) and iconic gestures (e.g., raising the fist to the ear to mean telephone). This is hypothesized to demonstrate the child's prelinguistic intentional communication and cooperation skills (Tomasello, Carpenter, and Liszkowski 2007).

Toward the beginning of the third year, the child starts to develop more complex grammatical constructs and skills. This is the case, for example, of the "verb islands" phenomenon (Tomasello 1992). Initially, children can use a variety of verbs and treat them as independent syntactic elements called "verb islands" (e.g., the child only uses very simple syntactic combinations of the same verb with different nouns of objects: "cut bread," "cut paper"). These intermediate syntactic constructions allow the child to subsequently develop more refined morphological and syntactic constructs, with more general verb islands combined with a richer set of prepositions. From the fourth year of age, the child gradually develops adultlike syntactic constructions such as simple transitives (agent-verb-patient, as in "John likes sweets") and locatives (agent-verb-patient-locative-location, as in "John puts sweets on table"; Tomasello and Brooks 1999). This gradually leads to the development of ever-more complex syntactic-morphologic constructions, more abstract and generalized grammatical categories known as word classes. These syntactic skills are accompanied by extended pragmatic and communicative skills, leading to refined narrative and discursive capabilities.

The constructivist view of language is highly consistent with the embodied and situated cognition theories (Pezzullo et al. 2013) and the relevant embodied robotics approach to the modeling of language learning (Cangelosi 2010, 2011). This embodied view stresses

Table 20.1
Typical timescale and major milestones of language development

Age (months)	Competence
0–6 months	Marginal babbling
6–9 months	Canonical (reduplicated) babbling
10–12 months	Intentional communication First gestures
12 months	Single words, holophrases Word-gesture combinations
18 months	Reorganization of phonological representations 50+ word lexicon size, vocabulary spurt Two-word combinations
24 months	Increasingly longer multiple-word sentences Verb islands
36+ months	Adultlike grammatical constructions Narrative skills

Source: Adapted from Cangelosi and Schlesinger 2015.

the fact that the body of the child, and its interaction with the environmental context, determines the type of representations, internal models, and cognitive strategies learned.

In cognitive robotics models, the embodied approach is linked to that of “symbol grounding” (Harnad 1990; Cangelosi 2010) and “grounded cognition” (Pezzulo et al. 2013). This refers to the capability of natural and artificial cognitive agents to acquire an intrinsic (autonomous) link between internal symbolic representations and referents in the external world or internal states. Cognitive robotics models implement the grounded learning of associations between words and the external and internal entities they refer to (objects, actions, internal states).

20.1.2 Levels of Analysis in Language Studies

In linguistics and psychology, a hierarchy of five levels of language analyses has been proposed: phonetic, lexical, semantic, syntactic, and pragmatic (see Cangelosi 2017). These levels are useful in cognitive robotics models because they identify the different aspects that need to be modeled and implemented to successfully achieve humanlike linguistic capabilities. For example, a robot, like a person, must be able to recognize language-specific sounds (phonetic level) to segment and identify the words (lexical level) and the grammatical structure of spoken utterances (syntactic level). This supports the understanding of the meaning of words and sentences (semantic level) and their contextualization within the interactive communication task (pragmatic level). These different levels of analysis should not, however, be considered separate modular components of language-processing models. In fact, all levels of language are strictly intertwined. For example, knowledge of the lexicon helps the lower-level recognition of phonemes and words. The pragmatic level of communication can also prime the recognition of the words and sentences that the hearer expects the speaker to choose to communicate the intended meaning.

Cognitive robotics models of language benefit from the field of natural language processing (NLP), which uses a set of computation linguistics methods for the different levels

of analysis and the representation of language. Numerous NLP methods and software tools have been proposed for phonetic analysis and automatic speech recognition (e.g., Markov models), for lexical and semantic analysis (e.g., WordNet), for parsing and syntactic analysis, and for pragmatics and communication (e.g., dialogue systems). This field has very recently gone through a significant revolution with the use of deep-learning models (cf. chapter 5). For example, deep neural networks are used for state-of-the-art speech recognition systems and parsing and word tagging (LeCun et al. 2015). These changes include the increasing use of end-to-end (a.k.a seq-to-seq—i.e., sequence-to-sequence) machine-learning models. These use deep neural networks that receive the raw input (e.g., sound wave or a word list) and, without specifically decomposing the linguistic processing into different levels of analysis or mechanisms, produce the desired output (e.g., translation of the input sentence into another language). In section 20.2 we will look at both NLP and the deep-learning models used in language systems for cognitive robots.

20.2 Robot Language Models

In robot language research, we can distinguish three main approaches to the design of language communication capabilities in robots (Cangelosi and Ogata 2019). The first directly models incremental, developmental phenomena on language acquisition. This is primarily based on developmental robotics approaches (chapter 3). Another approach is based on various NLP techniques, while the third focuses on the latest machine-learning approaches (chapter 9). The NLP approach typically combines off-the-shelf techniques and language-processing tools (e.g., ready-made lexicons and knowledge bases, parsers, automatic speech recognition, and speech synthesis software) to implement in the robot the ability to respond to linguistic instructions and to utter sentences to express a request. The language-learning approach, on the other hand, uses machine-learning methods (e.g., neural networks, Bayesian methods) to train the robot to acquire language skills. In practice, however, some NLP robotic approaches do use machine-learning methods (e.g., most of the current speech recognition systems are based on statistical learning and deep neural network methods), and some robot language-learning approaches partially rely on off-the-shelf NLP tools.

20.2.1 Developmental Robot Language Models

Developmental language-learning models are typically based on the developmental robotics approach (Cangelosi and Schlesinger 2015; see also chapter 3). As such, this approach puts a strong emphasis on constraining the robot's cognitive and linguistic architecture and behavioral and learning performance to known child psychology theories, data, and developmental principles. This permits the modeling of the developmental sequence of the qualitative and quantitative stages leading to the acquisition of adultlike sensorimotor, cognitive, and linguistic skills. Developmental robotics is also naturally suited to model embodied and situated cognition for the grounding of cognition (Pezzulo et al. 2013). Specifically, for the embodied bases of language learning, the use of robots that have to learn to name objects they see and name actions they perform constitutes an ideal way to model the grounding of symbols in sensorimotor knowledge and experience (Harnad 1990; Cangelosi 2011).

Some developmental robotics models focus on the acquisition and grounding of the first words. These models directly rely on child psychology studies on language acquisition in infants in the second year of age—that is, when the first words are acquired. One seminal developmental model is that of Morse et al. (2010, 2015), as it directly replicates child psychology data on embodied language acquisition via body posture interaction (Samuelson et al. 2011). In Samuelson et al.'s (2011) child psychology study, the infant repeatedly experiences two new objects (the target and the foil) in different locations (left/right), requiring a postural change to attend to the object. Subsequently, the child hears the object name “modi” while attending to a foil object that has been placed in the location normally associated with the target object. When the infant is asked, “Where is the modi?,” they select the target object—that is, the object normally associated with the posture and spatial location they were attending to, rather than the actual object they were looking at when they heard the name. This means that infants rely on memory for their own posture and the related object location to associate objects and their names.

Morse et al. (2015) have proposed an embodied model of this phenomenon with the iCub humanoid robot, replicating the original experiments by Samuelson et al. and further exploring how this spatial component can be achieved via the robot's physical interaction with objects and locations. The model is an implementation of the epigenetic robotics architecture (Morse et al. 2010), a developmental robotics cognitive architecture specifically designed for studying embodied language learning. The core of such an architecture consists of three self-organizing maps with Hebbian connections between their units (figure 20.1). The first (visual) map is used to represent in a topological way the similarity of preprocessed visual information (e.g., color and/or shape) implemented as input of a spectrogram of the color of each object in view. The second (body) map is driven by postural information (the current motor encoder values of the eyes, head, and torso of the robot). The final (word) map responds to each word encountered (preprocessed by a standard NLP speech recognition system). The visual color map and the word map are both fully connected to the body posture map, with connection weights adjusted by a normalized positive and negative Hebbian learning rule.

In one version of the experiment, the target object (a red ball) is placed to the left of the iCub. The robot looks at the target for approximately ten seconds before the target object is removed, and the foil object is placed to the right of the iCub, which again orients for approximately ten seconds. This procedure is repeated four times. In the fifth presentation cycle, the foil object is placed in the position normally associated with the target object, and the word “modi” is spoken. The original placements of each object are repeated one final time, and then both objects are positioned in new locations to test the robot by stating, “Find the modi.” The robot then orients and reaches for one of the objects. Various versions of the experiment were carried out, each repeated twenty times (with all prelearning weights randomly initialized). Morse et al. (2015) conducted an additional experiment following the same procedure outlined above but with the addition of another spatial dimension of the robot's posture (from sitting to standing) for the naming event only at the fifth presentation cycle. As a result of this change, the naming event occurs in a posture that has not been previously associated with either the target or the foil object. Thus, testing the interference between previously experienced objects and that posture causes the iCub to select the foil object (the object it was observing when it first heard the name). This

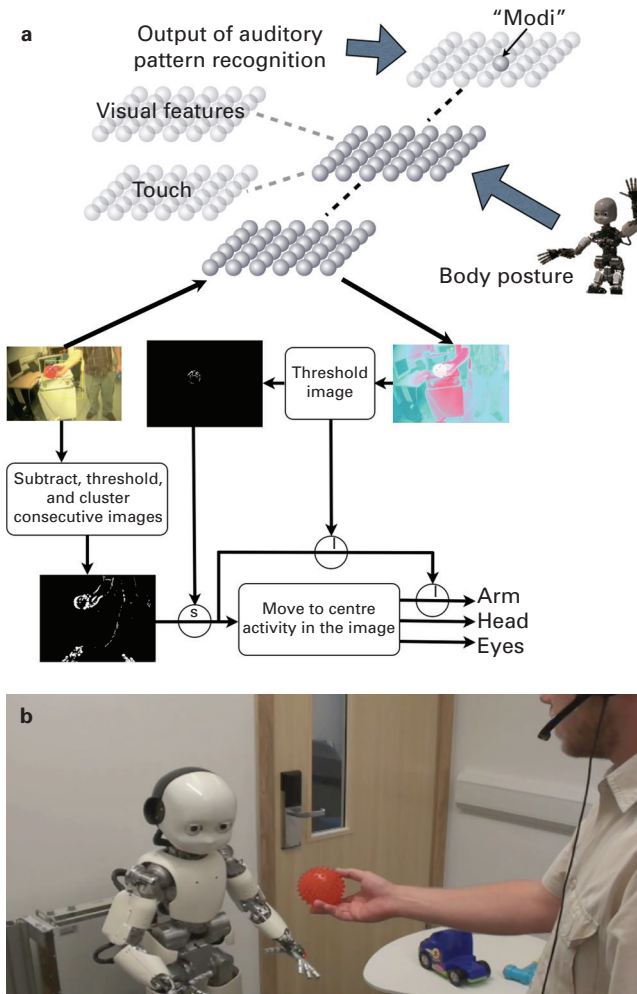


Figure 20.1

Setup for word-learning experiments (a) and cognitive architecture (b) in Morse et al. (2015).

result was also replicated in new child experiments (Morse et al. 2015). Overall, this model shows that infants, like robots, use the memory of postures as a way to organize their learning task. If two different postures are used at this early stage of development, they are used by the robot to separate different cognitive tasks.

An extended version of this model has already been used to replicate a range of other language acquisition phenomena (Morse and Cangelosi 2017; Cangelosi and Schlesinger 2018). For example, Twomey et al. (2016) used the ERA architecture to model mutual exclusivity—that is, the developmental phenomenon in which a child can learn the name of a new object if they hear a new label and are presented with an unseen (unlabeled) object among other objects with a known label. Other developmental language models have looked at the learning of both object and action labels, moving toward the first examples of syntax learning. For example, Tikhonoff et al. (2011) proposed a simulation

model of the iCub robot in the development of a lexicon based on both names of objects and of actions and their basic combinations to understand simple commands such as “pick_up blue_ball.”

A few developmental robotics models have focused on grammar development—for example, modeling the emergence of semantic compositionality for syntactic compositionality for multiple word combinations and generalizations (Sugita and Tani 2005; Tuci et al. 2011; Zhong et al. 2019). For example, the robot model by Sugita and Tani (2005) investigated the emergence of compositional meanings and lexicons with no a priori knowledge of any lexical or formal syntactic representations. The environment consisted of three colored objects (red, blue, and green) in three different locations (a red object on the left-hand side of the robot’s field of view, a blue object in the middle, and a green object on the right). The robot could respond with nine possible behaviors based on the combination of three actions (POINT, PUSH, HIT) with the three objects (RED, BLUE, GREEN) always in the same locations (LEFT, CENTER, RIGHT). The robot learning architecture was a parametric bias recurrent neural network (PBRNN), which is capable of learning a set of parametric bias units able to represent action sequences via language-like symbols. The robot experiments were divided into two stages: training and generalization. In the training phase, the robot acquired associations between given sample training sentences and corresponding behavioral sequences. In the testing phase, the robot’s ability to generate the correct behavior by recognizing the sentences used during training and, above all, novel combinations of words was tested. A subset of fourteen object/action/location combinations was used during training, with four left for the generalization test. After the successful training stage, in the generalization test phase the four remaining novel sentences were given to the robot: “Point green,” “point right,” “push red,” and “push left.” Behavioral results showed whether the linguistic module had acquired the underlying compositional syntax correctly. The robot could generate grammatically correct sentences and understand them by giving a behavioral demonstration of the generalized actions. Detailed analyses of the robot’s neural representations supporting the verb-noun compositional knowledge showed a separated substructure for the verbs and nouns. In particular, the congruence in the substructures for verbs and nouns indicated that the combinatorial semantic/syntactic structure was successfully extracted by the robot’s neural network.

Yamashita and Tani (2008) proposed an extension of this work using the multiple-timescale recurrent neural network (MTRNN) for compositional action and language learning experiments. Zhong et al. (2019) further extended the MTRNN architecture to control the compositional learning and generalization of nine actions on nine objects for verb-noun learning in the iCub robot.

Developmental learning models have also been proposed to investigate the acquisition of abstract concepts and words in robots, including words referring to general-purpose motor actions such as “use” and “make” and number and counting words (Cangelosi and Stramandinoli 2018). To model the grounding and embodied bases of abstract word learning in cognitive robots, one study looked at abstract action verbs such as “to use,” which can be applied to different motor contexts (e.g., “use a hammer” or “use a pen”) with no common motor program. The developmental robotics model of Stramandinoli et al. (2017) exploits the hierarchical recursive structures of both the linguistic and the motor system to integrate simple motor primitives and concrete words to create the semantic referents

of abstract action words that do not have a direct mapping to the sensorimotor world. An iCub robot is first trained to recognize a set of tools of different colors, sizes, and shapes (e.g., knife, hammer, brush) and to perform object-related actions (e.g., cut, hit, paint). Subsequently, the robot is taught to name these objects and actions (e.g., “cut with knife”). Finally, the robot is taught the abstract motor words of “use” and “make” by combining these new action words with the appropriate tool name (e.g., “use knife”). The experiments investigated the effects of using different combinations of the three input modalities (i.e., vision, language, and proprioception). For example, incompatible condition tests between the perceptual and linguistic input showed that the robot ignored the linguistic command by executing the actions elicited by the seen objects. Hence, the knowledge associated with objects relies not only on the objects’ perceptual features but also on the actions that can be performed on them (i.e., affordances). Further simulation experiments showed that the acquisition of concepts related to abstract action words (e.g., “use knife”) requires the reactivation of similar internal representations of the network activated during the acquisition of the concrete concepts (e.g., “cut with knife”) contained in the linguistic sequences used for the grounding of abstract action words (e.g., “use knife” is “cut with knife”). This finding suggests that the semantic representation of abstract action words requires the recall and reuse of sensorimotor representational capabilities (i.e., embodied understanding of abstract language). Indeed, neurophysiological evidence of the modulation of the motor system during the comprehension of both concrete and abstract language exists to support this finding.

Finally, developmental models with humanoid robots have also been used to model abstract concepts and the representation of the underlying knowledge of numbers. Number cognition is another key example of the contribution of embodied cognition in the acquisition of abstract, symbol-like manipulation capabilities. Various embodied strategies, such as pointing and counting gestures, object touching, and finger counting, have been shown to facilitate the development of number cognition skills (e.g., Alibali and DiRusso 1999; Moeller et al. 2011). Given the implicit embodied nature of humanoid robots, some recent models have specifically looked at the modeling of the acquisition of number concepts and words via embodied strategies such as gestures (Ruciński et al. 2012) and finger counting (De La Cruz et al. 2014; Pecyna et al. 2020). For example, a developmental robotics model was used specifically to explore whether finger counting and the association of number words to each finger could bootstrap the representation of numbers in a cognitive robot. This study used a recurrent artificial neural network to model the learning of associations between (motor) finger counting, (visual) object counting, and (auditory) number word and sequence learning. In particular, this study manipulated the coupling between different modalities, such as with the comparison of the Auditory-Only condition, when the robot solely learns to hear and repeat the sequence of number words (“one,” “two,” . . . up to “ten”), with the Finger+Auditory condition, when the robot simultaneously learns the sequence of acoustic number words and moving fingers.

The results showed that learning the number word sequences together with finger sequencing (Finger+Auditory condition) helps to quickly build the initial representation of numbers in the robot. Robots who only learn the auditor sequences (Auditory-Only condition) achieve the worst performances. Moreover, the neural network’s internal representations of these two conditions resulted in qualitatively different patterns of similarity

in the representation between numbers. Only after the Finger + Auditory sequence learning did the network represent the relative distance between numbers, which corresponded to the quantitative difference between numbers. In Finger+Auditory-trained robots, the cluster analysis diagram of the hidden layer's activation showed that the representation for the word "one" was adjacent to that of "two" and increased differently (distant) from the higher numbers. However, in the auditory-only condition, there was no correspondence between the cluster diagram similarity distance and the numerical distance.

This finger-counting model has recently been extended by Pecyna et al. (2020) to model numerosity estimation and by Di Nuovo and McClelland (2019), who combined developmental robotics and deep-learning methods to show that proprioceptive information from robot hands improves accuracy in the recognition of spoken digits. See chapter 22 for an extended discussion of abstract and number word learning.

20.2.2 NLP-Based Robot Language Models

NLP methods have been used for two different types of robot language models. In the conversational approach, the robot uses NLP tools primarily to engage in a linguistic conversation with a human user for social companionship, entertainment, or information-gathering tasks, with no actual motor tasks to perform (no language grounding required). In human-robot interaction models, robots use language primarily to respond to instructions to perform a physical action.

Conversational robots have their origins in conversational agents and chatterbots, such as the very first conversational agent developed called ELIZA (Weizenbaum 1966). More recent conversational agents are often based on animated virtual 3D characters, such as A.L.I.C.E. (Wallace 2009). Conversational agents embodied in physical robots include work with the android robot ERICA (ERato Intelligent Conversational Android; Ishiguro 2016), the Robot-ERA system for supporting older people in independent living (Di Nuovo et al. 2018), and museum/station guides and robot tutors for children (Shiomi et al. 2008; Belpaeme et al. 2018). These conversational robots use a variety of NLP tools for speech recognition, parsing, and dialogue systems.

Many NLP-based robot language systems are designed with the primary function of following a user's instructions and selecting the appropriate motor behavior. These applications typically cover object manipulation tasks (e.g., "pick up blue ball," "clean the table") and navigation scenarios (e.g., "go to the exit," "take me to the restroom"; Mavridis 2015). The use of speech for language instruction understanding requires a tight coupling (grounding) of the robot's visual and motor repertoire with its language processing and knowledge representation methods. In NLP-based approaches, this link is typically predefined by the designer. There is no autonomous grounding of the robot's words via situated learning, as the robot can only use a set of "meanings" defined by the programmer. For example, Aloimonos and Pastra developed a language and action representation formalism, called PRAXICON, for action and language knowledge representation of object manipulation tasks (Pastra and Aloimonos 2012; Pastra 2008). It uses a goal-based representation of actions employing a multimodal semantic network-type representation that is directly inspired by linguistic methods, such as the mapping of a minimalist grammar of language into a minimalist grammar of action representation. PRAXICON was tested on the Baxter robot capable of learning to cook from "watching" videos available on the Web (Yang et al. 2015).

Nonverbal communication capabilities have also been proposed to complement and enhance a robot's linguistic production and communicative expressivity. For example, Csapo et al. (2012) complemented speech production with nonverbal strategies such as face tracking, nodding, gesturing, proximity detection, and interruptions. Mutlu et al. (2012) modeled humanlike gaze mechanisms to help robots signal different interaction roles to the human interlocutor to manage turn exchanges and the dynamics of the conversation.

20.2.3 Machine-Learning Robot Language Models

Multimodal integration, which directly concerns the field of language learning for connecting speech, vision, and action, has long been a difficult problem in robotics. For example, the crossmodal complementation of information loss or the application of crossmodal memory search for behavior generation problems have not been thoroughly studied. Second, literature discussions on how to fuse multimodal information to achieve stable environmental awareness have not reached a comprehensive consensus. In robotics, the sensory input acquired from different sources is still typically processed using a dedicated feature extraction mechanism (Murphy 2019). Third, multimodal synchronization modeling as a means for implementing the sensorimotor prediction of robot applications has not been adequately studied. Several studies so far have proposed a computational model that develops synchronization of behavioral effects in a developmental way toward an understanding of interaction (Kuriyama et al. 2010; Ogino et al. 2006). However, most casual models are expressed using a limited number of modalities and in many cases focus only on vision and behavior.

In recent years, different types of graphic models of multimodal classification have been reported. Lallee and Dominey (2013) proposed a multimodal convergence map based on a self-organizing map (SOM) that integrates visual-motor and language modality. Sinapov and Stoytchev (2011) developed a graph-based model that enables robots to recognize untrained objects based on their similarity to trained objects. They also let the robot take ten different actions to collect visual, auditory, and tactile data; explore one hundred objects; and categorize twenty objects with supervised learning (Sinapov et al. 2014). Ivaldi et al. (2013) developed a robot that can learn object categories by active sensing. Nakamura et al. (2009, 2015) proposed studies on multimodal classification using multi-lateral latent Dirichlet allocation (MLDA) and its extension. They developed robotic systems that can obtain visual, sound, and tactile information by handling objects. The robot grasps an object several times and shakes the object to acquire sound information. By applying the MLDA, they showed that robots can classify many objects into categories, which is similar to human classification results (Nakamura et al. 2009). Araki et al. (2011) developed an MLDA online and conducted experiments on completely autonomous multimodal category acquisition in the home environment.

Notwithstanding the above multimodal machine-learning examples, there has been little research on scalable learning frameworks for handling a large amount of sensorimotor data of a high dimension. The latest robots are equipped with state-of-the-art sensor devices such as high-resolution image sensors, distance sensors, and multichannel microphones as the demand for perception accuracy with respect to the surrounding environment increases (Kaneko et al. 2008; Sakagami et al. 2002). Thus, a remarkable improvement in the amount of sensorimotor information available has been achieved. However, due to the

scalability limitations of conventional machine-learning algorithms, few computational models achieve robust behavior control and environmental recognition by fusing multimodal perceptual inputs into a single representation. To overcome the problem of the scalability limitation, deep-learning approaches such as deep neural networks (DNNs), used as perceptual feature extraction and multimodal integration learning mechanisms, have attracted the attention of the robotics and machine-learning community in recent years. One of the main advantages of applying a DNN is the ability to self-organize highly generalized sensory functions from large-scale raw data. For example, DNNs have been successfully applied to unsupervised feature learning for a single modality such as text, images, and voice. The same approach has also been applied to the learning of integrated representation among multiple modalities, resulting in a significant improvement in speech recognition performance. In another context using unsupervised learning, Le (2013) showed that DNNs with large-scale data can automatically construct high-level features from image data. Connecting acquired representation by neural networks and multimodal classification is an important research field (Bengio et al. 2013). However, the application of DNNs for more dynamic information such as robot motion and language has just begun to be considered.

Multimodal integration based on DNNs is generally accomplished by two approaches. First, in the feature extraction method, feature vectors from some plural modalities are transformed to acquire an integrated feature vector. For example, Ngiam et al. (2011) utilized a DNN that extracts directly integrated expressions from multimodal signal input by compressing the input dimension. Huang and Kingsbury (2013) used deep belief networks (DBNs) for audiovisual speech recognition tasks by combining intermediate-level features learned by a DBN of a single modality. However, these methods have difficulty explicitly and adaptively selecting their respective information gains in response to dynamic changes in the reliability of multimodal information sources. Alternatively, in the fusion method the outputs of the unimodal classifiers are merged to determine the final classification. Unlike the feature extraction approaches, the fusion methods can improve robustness by incorporating the stream reliability associated with multiple information sources as a measure of the information gain of the recognition model.

Specifically for robot language models, Noda et al. (2015) proposed a speech recognition model that uses a DNN both for noise reduction of speech features and for using visual information in a complementary style. The perception features acquired from the audio signal and the corresponding mouth region image are then integrated. Two kinds of DNNs, a deep denoising autoencoder (DDA) and a convolution neural network (CNN), are used for the feature extraction of audio information and visual information, respectively. In addition, the multistream hidden Markov model (MSHMM) is applied to integrate the two perceptual features acquired from the speech signal and the mouth region image. They show that the CNN outputs higher recognition rates than the visual features extracted by PCA (principal component analysis), and the effect of the different image resolutions is not prominent. The word recognition rate, visual features acquired by the CNN, is approximately 22.5 percent.

The DNN language and multimodal integration models provide intuitive and direct ways to accomplish temporal sequence recognition tasks. The focus of the task is to “recognize” by symbolizing the raw sensory signal. However, since recognition methods using

probabilistic models specialize in obtaining symbolic representation from the raw signal, they are not suitable for sensorimotor coordination tasks, such as robot behavior generation. Therefore, this approach needs to design external mechanisms to generate behaviors corresponding to the recognized state. To address this, Heinrich et al. (2015) utilized multiple timescale recurrent neural networks (MTRNN) to integrate visual, auditory, and motor information.

Noda et al. (2014) also proposed a multimodal temporal sequence integration learning framework using a DNN for multimodal time series integrated learning, as well as feature extraction by dimensional compression. They showed the framework with multiple DNNs as a crossmodal memory retriever and as a temporal sequence predictor. Specifically, they integrated image, sound signal, and motor modalities with multiple deep autoencoders (DAs). The learning experiments were conducted on six types of object manipulations by the humanoid robot NAO, generated by direct teaching. The data of high dimension, such as images and sound signals, are compressed to thirty dimensions by the DA. The image and sound data obtained from this process and the motor command obtained from the robot are integrated using a DA instead of an HMM. The data were extracted within a sliding time window of thirty steps. Results showed that this model self-organizes not only the sensory features but also the motion patterns from the time series of sensorimotor data corresponding to the plural robot motions. The principal component analysis of the acquired internal representation showed that each motion does not correspond to the motion cluster designed by the human teacher. Some motions have multiple clusters reflected by the characteristics of the learning condition. Some motions overlap the other motions, thereby associating with each other. Thus, the real world, the body structure, and the learning model self-organize the expressions of behaviors coupled with recognition. They realized a cross-modal memory association by using this internal representation. For example, robot motion is generated from images and sound data; the visual image (movie) is produced from body motion or sound. This demonstrates a significant advantage of using DNN multimodal learning to generate expressions of a very large dimension.

DNNs have also been used to extend developmental models of language learning, integrating recurrent neural networks, such as long short-term memory (LSTM), with simultaneous action and language processing. For example, Antunes et al. (2019) used a bidirectional multiple timescales LSTM for the grounding of actions and verbs without explicitly learning an intermediate representation. The model self-organizes such representations at the level of a slowly varying latent layer connecting the language and the action route (figure 20.2). The model is also trained in a bidirectional way, learning how to produce a sentence from a certain action sequence input and, simultaneously, how to generate an action sequence given a sentence as input. This network was evaluated on motor actions performed by an iCub robot and their corresponding letter-based description. Yamada et al. (2017) also used recurrent DNNs to train a robot to translate sentences that included logic words, such as “not,” “and,” and “or,” into robot actions. The model analysis showed that referential words are merged with visual information and the robot’s own current state, while logical words are represented by the model in accordance with their functions as logical operators.




Figure 20.2

A bidirectional LSTM for action and language learning. *Source:* Adapted from Antunes et al. 2019.

20.3 Conclusion

This chapter summarizes the cognitively inspired approaches to the design of language learning and language grounding and processing capabilities in robots. Developmental language-learning models have been able to replicate humanlike developmental trajectories in the early acquisition of words and simple grammatical structures. They also exploit embodied strategies, such as posture bias and finger-counting skills, in learning and grounding concrete and abstract words. However, the level of complexity of the robot's language repertoire is limited to small lexicons. NLP-based models, on the other hand, have been widely used to handle dialogue with conversational agents and complex lexicons. However, in these models the robot is not able to autonomously ground the words it uses for sensorimotor knowledge, and it must rely on the hand coding of the word-meaning mappings defined by the system designer.

An important development in robot language research is the very recent progress on learning methods for language and multimodal information based on machine-learning models. However, on its own, DNN cannot address the whole problem of robot language grounding. For example, deep learning takes a batch-learning and a supervised-learning approach, and generally, it cannot work online. It acquires representations approximating the given input data, and it cannot easily define novel symbols (and meanings) about the world, as humans do with language generativity. It is also important to acknowledge that although DNNs can match human performance in some particular data-processing tasks,

they do have significant limitations. The most critical issue with DNNs for robot language models is that it is extremely challenging to understand a DNN's internal mechanism. Even when high performance is achieved, it is difficult to identify the cause when a mistake occurs. This is a serious problem in the behavior learning of real-world systems such as an interactive robots or automatic driving cars. In DNNs, the internal representation is embedded not only in its large structure but also in its small structure. These mechanisms enable DNNs to self-organize very large and complicated structures of data and to show high performance rates. However, simple statistical analysis and modeling are not directly effective for explaining the mechanism of deep learning. Thus, a mathematical understanding of the DNN as a multidimensional complex system—that is, a dynamic system—is an important area for future work that will have significant implications for the use of deep learning in robot language models.

Finally, an important direction for future research is to focus on a developmental approach, where symbol acquisition emerges from the incremental interaction between the robot, the human user, and their environment. This requires the long-term and open-ended development of a human-robot interaction and communications system that allows a developmental learning robot to bootstrap its multimodal, grounded language-learning skills and repertoire.

Additional Reading and Resources

- An extensive position paper proposing a developmental robotics approach to communication and language integration: Cangelosi, Angelo, Giorgio Metta, Gerhard Sagerer, Stefano Nolfi, Chrystopher Nehaniv, Kerstin Fischer, Jun Tani, et al. 2010. "Integration of Action and Language Knowledge: A Roadmap for Developmental Robotics." *IEEE Transactions on Autonomous Mental Development* 2 (3): 167–195.
- A comprehensive paper on the symbol-emergence approach to language development modeling: Taniguchi, Tadahiro, Takayuki Nagai, Tomoaki Nakamura, Naoto Iwahashi, Tetsuya Ogata, and Hideki Asoh. 2016. "Symbol Emergence in Robotics: A Survey." *Advanced Robotics* 30 (11–12): 706–728.
- A recent extensive review of language and speech models for humanoid robotics: Cangelosi, Angelo, and Tetsuya Ogata. 2019. "Speech and Language in Humanoid Robots." In *Humanoid Robotics: A Reference*, edited by P. Vadakkepat and A. Goswami. Berlin: Springer.

References

- Alibali, Martha Wagner, and Alyssa A. DiRusso. 1999. "The Function of Gesture in Learning to Count: More than Keeping Track." *Cognitive Development* 14 (1): 37–56.
- Antunes, Alexandre, Alban Lafflaquière, Tetsuya Ogata, and Angelo Cangelosi. 2019. "A Bi-directional Multiple Timescales LSTM Model for Grounding of Actions and Verbs." In *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2614–2621. New York: IEEE.
- Araki, Takaya, Tomoaki Nakamura, Takayuki Nagai, Kotaro Funakoshi, Mikio Nakano, and Naoto Iwahashi. 2011. "Autonomous Acquisition of Multimodal Information for Online Object Concept Formation by Robots." In *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1540–1547. New York: IEEE.
- Belpaeme, Tony, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. 2018. "Social Robots for Education: A Review." *Science Robotics* 3 (21).

- Bengio, Yoshua, Aaron Courville, and Pascal Vincent. 2013. "Representation Learning: A Review and New Perspectives." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8): 1798–1828.
- Cangelosi, Angelo. 2010. "Grounding Language in Action and Perception: From Cognitive Agents to Humanoid Robots." *Physics of Life Reviews* 7 (2): 139–151.
- Cangelosi, Angelo. 2011. "Solutions and Open Challenges for the Symbol Grounding Problem." *International Journal of Signs and Semiotic Systems* 1 (1): 49–54.
- Cangelosi, Angelo. 2017. "Language Processing." In *From Neuron to Cognition via Computational Neuroscience*, edited by M. Arbib and J. Bonaiuto, 693–718. Cambridge, MA: MIT Press.
- Cangelosi, Angelo, and Tetsuya Ogata. 2019. "Speech and Language in Humanoid Robots." In *Humanoid Robotics: A Reference*, edited by P. Vadakkepat and A. Goswami, 2261–2292. Berlin: Springer.
- Cangelosi, Angelo, and Matthew Schlesinger. 2015. *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Cangelosi, Angelo, and Matthew Schlesinger. 2018. "From Babies to Robots: The Contribution of Developmental Robotics to Developmental Psychology." *Child Development Perspectives* 12 (3): 183–188.
- Cangelosi, Angelo, and Francesca Stramandinoli. 2018. "A Review of Abstract Concept Learning in Embodied Agents and Robots." *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170131.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Csapo, Adam, Emer Gilmartin, Jonathan Grizou, Jingguang Han, Raveesh Meena, Dimitra Anastasiou, Kristiina Jokinen, and Graham Wilcock. 2012. "Multimodal Conversational Interaction with a Humanoid Robot." In *2012 IEEE 3rd International Conference on Cognitive Infocommunications*, 667–672. New York: IEEE.
- De La Cruz, Vivian Milagros, Alessandro Di Nuovo, Santo Di Nuovo, and Angelo Cangelosi. 2014. "Making Fingers and Words Count in a Cognitive Robot." *Frontiers in Behavioral Neuroscience* 8:13.
- Di Nuovo, Alessandro, F. Broz, N. Wang, T. Belpaeme, A. Cangelosi, R. Jones, R. Esposito, F. Cavallo, and P. Dario. 2018. "The Multi-modal Interface of Robot-Era Multi-robot Services Tailored for the Elderly." *Intelligent Service Robotics* 11 (1): 109–126.
- Di Nuovo, Alessandro, and Jay L. McClelland. 2019. "Developing the Knowledge of Number Digits in a Child-Like Robot." *Nature Machine Intelligence* 1 (12): 594–605.
- Goldberg, Adele E. 2006. *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.
- Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D* 42:335–346.
- Heinrich, Stefan, Sven Magg, and Stefan Wermter. 2015. "Analysing the Multiple Timescale Recurrent Neural Network for Embodied Language Understanding." In *Artificial Neural Networks*, edited by P. Koprinkova-Hristova, V. Mladenov, and N. K. Kasabov, 149–174. Cham, Switzerland: Springer.
- Hoff, Erika. 2013. *Language Development*. Boston: Cengage Learning.
- Huang, Jing, and Brian Kingsbury. 2013. "Audio-Visual Deep Learning for Noise Robust Speech Recognition." In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 7596–7599. New York: IEEE.
- Ishiguro, Hiroshi. 2016. "Android Science." In *Cognitive Neuroscience Robotics A*, edited by M. Kasaki, H. Ishiguro, M. Asada, M. Osaka, and T. Fujikado, 193–234. Tokyo: Springer.
- Ivaldi, Serena, Natalia Lyubova, Alain Droniou, Vincent Padois, David Filliat, Pierre-Yves Oudeyer, and Olivier Sigaud. 2013. "Object Learning through Active Exploration." *IEEE Transactions on Autonomous Mental Development* 6 (1): 56–72.
- Kaneko, Kenji, Kensuke Harada, Fumio Kanehiro, Go Miyamori, and Kazuhiko Akachi. 2008. "Humanoid Robot HRP-3." In *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2471–2478. New York: IEEE.
- Kuriyama, Takatsugu, Takashi Shibuya, Tatsuya Harada, and Yasuo Kuniyoshi. 2010. "Learning Interaction Rules through Compression of Sensori-motor Causality Space." In *Proceedings of The Tenth International Conference on Epigenetic Robotics (Epirob10)*, 57–64. Lund University Cognitive Studies, 149.
- Lalée, Stéphane, and Peter Ford Dominey. 2013. "Multi-modal Convergence Maps: From Body Schema and Self-Representation to Mental Imagery." *Adaptive Behavior* 21 (4): 274–285.
- Le, Quoc V. 2013. "Building High-Level Features Using Large Scale Unsupervised Learning." In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 8595–8598. New York: IEEE.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–444.
- Liu, Rui, and Xiaoli Zhang. 2019. "A Review of Methodologies for Natural-Language-Facilitated Human–Robot Cooperation." *International Journal of Advanced Robotic Systems* 16 (3): 1729881419851402.

- MacWhinney, B. 1998. "Models of the Emergence of Language." *Annual Review of Psychology* 49:199–227.
- Mavridis, Nikolaos. 2015. "A Review of Verbal and Non-verbal Human–Robot Interactive Communication." *Robotics and Autonomous Systems* 63:22–35.
- Moeller, Korbinian, Laura Martignon, Silvia Wesselowski, Joachim Engel, and Hans-Christoph Nuerk. 2011. "Effects of Finger Counting on Numerical Development—the Opposing Views of Neurocognition and Mathematics Education." *Frontiers in Psychology* 2:328.
- Morse, Anthony F., Viridian L. Benitez, Tony Belpaeme, Angelo Cangelosi, and Linda B. Smith. 2015. "Posture Affects How Robots and Infants Map Words to Objects." *PLoS One* 10 (3): e0116012.
- Morse, Anthony F., and Angelo Cangelosi. 2017. "Why Are There Developmental Stages in Language Learning? A Developmental Robotics Model of Language Development." *Cognitive Science* 41:32–51.
- Morse, Anthony F., Joachim de Greeff, Tony Belpaeme, and Angelo Cangelosi. 2010. "Epigenetic Robotics Architecture (ERA)." *IEEE Transactions on Autonomous Mental Development* 2 (4): 325–339.
- Murphy, Robin R. 2019. *Introduction to AI Robotics*. Cambridge, MA: MIT Press.
- Mutlu, Bilge, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. 2012. "Conversational Gaze Mechanisms for Humanlike Robots." *ACM Transactions on Interactive Intelligent Systems* 1 (2): 1–33.
- Nakamura, Tomoaki, Yoshiki Ando, Takayuki Nagai, and Masahide Kaneko. 2015. "Concept Formation by Robots Using an Infinite Mixture of Models." In *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4593–4599. New York: IEEE.
- Nakamura, Tomoaki, Takayuki Nagai, and Naoto Iwahashi. 2009. "Grounding of Word Meanings in Multimodal Concepts Using LDA." In *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3943–3948. New York: IEEE.
- Ngiam, Jiquan, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y. Ng. 2011. "Multimodal Deep Learning." In *ICML '11: Proceedings of the 28th International Conference on Machine Learning*. Madison, WI: Omnipress.
- Noda, Kuniaki, Hiroaki Arie, Yuki Suga, and Tetsuya Ogata. 2014. "Multimodal Integration Learning of Robot Behavior Using Deep Neural Networks." *Robotics and Autonomous Systems* 62 (6): 721–736.
- Noda, Kuniaki, Yuki Yamaguchi, Kazuhiro Nakadai, Hiroshi G. Okuno, and Tetsuya Ogata. 2015. "Audio-Visual Speech Recognition Using Deep Learning." *Applied Intelligence* 42 (4): 722–737.
- Ogino, Masaki, Hideki Toichi, Yuichiro Yoshikawa, and Minoru Asada. 2006. "Interaction Rule Learning with a Human Partner Based on an Imitation Faculty with a Simple Visuo-motor Mapping." *Robotics and Autonomous Systems* 54 (5): 414–418.
- Pastra, Katerina. 2008. "PRAXICON: The Development of a Grounding Resource." In *Proceedings of the International Workshop on Human-Computer Conversation*, Bellagio, Italy.
- Pastra, Katerina, and Yiannis Aloimonos. 2012. "The Minimalist Grammar of Action." *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1585): 103–117.
- Pecyna, Leszek, Angelo Cangelosi, and Alessandro Di Nuovo. 2020. "A Robot That Counts Like a Child: A Developmental Model of Counting and Pointing." *Psychological Research*. <https://doi.org/10.1007/s00426-020-01428-8>.
- Pezzulo, Giovanni, Lawrence W. Barsalou, Angelo Cangelosi, Martin H. Fischer, Ken Mcrae, and Michael Spivey. 2013. "Computational Grounded Cognition: A New Alliance between Grounded Cognition and Computational Modeling." *Frontiers in Psychology* 3:612.
- Ruciński, Marek, Angelo Cangelosi, and Tony Belpaeme. 2012. "Robotic Model of the Contribution of Gesture to Learning to Count." In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics*, 1–6. New York: IEEE.
- Sakagami, Yoshiaki, Ryujin Watanabe, Chiaki Aoyama, Shinichi Matsunaga, Nobuo Higaki, and Kikuo Fujimura. 2002. "The Intelligent ASIMO: System Overview and Integration." In Vol. 3, *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2478–2483. New York: IEEE.
- Samuelson, Larissa K., Linda B. Smith, Lynn K. Perry, and John P. Spencer. 2011. "Grounding Word Learning in Space." *PLoS One* 6 (12): e28095.
- Shiomi, Masahiro, Daisuke Sakamoto, Takayuki Kanda, Carlos Toshinori Ishi, Hiroshi Ishiguro, and Norihiro Hagita. 2008. "A Semi-autonomous Communication Robot—a Field Trial at a Train Station." In *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction*, 303–310. New York: IEEE.
- Sinapov, Jivko, Connor Schenck, Kerrick Staley, Vladimir Sukhoy, and Alexander Stoytchev. 2014. "Grounding Semantic Categories in Behavioral Interactions: Experiments with 100 Objects." *Robotics and Autonomous Systems* 62 (5): 632–645.

- Sinapov, Jivko, and Alexander Stoytchev. 2011. "Object Category Recognition by a Humanoid Robot Using Behavior-Grounded Relational Learning." In *2011 IEEE International Conference on Robotics and Automation*, 184–190. New York: IEEE.
- Stramandinoli, Francesca, Davide Marocco, and Angelo Cangelosi. 2017. "Making Sense of Words: A Robotic Model for Language Abstraction." *Autonomous Robots* 41 (2): 367–383.
- Sugita, Yuuya, and Jun Tani. 2005. "Learning Semantic Combinatoriality from the Interaction between Linguistic and Behavioral Processes." *Adaptive Behavior* 13 (1): 33–52.
- Tikhanoff, V., A. Cangelosi, and G. Metta. 2011. "Language Understanding in Humanoid Robots: iCub Simulation Experiments." *IEEE Transactions on Autonomous Mental Development* 3 (1): 17–29.
- Tomasello, Michael. 1992. *First Verbs: A Case Study of Early Grammatical Development*. Cambridge: Cambridge University Press.
- Tomasello, Michael. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.
- Tomasello, Michael, and Patricia J. Brooks. 1999. "Early Syntactic Development: a Construction Grammar Approach." In *Development of Language*, edited by M. Barrett, 161–190. London: Psychology Press.
- Tomasello, Michael, Malinda Carpenter, and Ulf Liszkowski. 2007. "A New Look at Infant Pointing." *Child Development* 78 (3): 705–722.
- Tuci, Elio, Tomassino Ferrauto, Arne Zeschel, Gianluca Massera, and Stefano Nolfi. 2011. "An Experiment on Behavior Generalization and the Emergence of Linguistic Compositionality in Evolving Robots." *IEEE Transactions on Autonomous Mental Development* 3 (2): 176–189.
- Twomey, Katherine E., Anthony F. Morse, Angelo Cangelosi, and Jessica S. Horst. 2016. "Children's Referent Selection and Word Learning: Insights from a Developmental Robotic System." *Interaction Studies* 17 (1): 93–119.
- Wallace, Richard S. 2009. "The Anatomy of ALICE." In *Parsing the Turing Test*, edited by R. Epstein, G. Roberts, and G. Beber, 181–210. Dordrecht: Springer.
- Weizenbaum, Joseph. 1966. "ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine." *Communications of the ACM* 9 (1): 36–45.
- Yamada, Tatsuro, Shingo Murata, Hiroaki Arie, and Tetsuya Ogata. 2017. "Representation Learning of Logic Words by an RNN: From Word Sequences to Robot Actions." *Frontiers in Neurobotics* 11:70.
- Yamashita, Yuichi, and Jun Tani. 2008. "Emergence of Functional Hierarchy in a Multiple Timescale Neural Network Model: A Humanoid Robot Experiment." *PLoS Computational Biology* 4 (11): e1000220.
- Yang, Yezhou, Yi Li, Cornelia Fermuller, and Yiannis Aloimonos. 2015. "Robot Learning Manipulation Action Plans by 'Watching' Unconstrained Videos from the World Wide Web." In *Twenty-Ninth AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Zhong, Junpei, Martin Peniak, Jun Tani, Tetsuya Ogata, and Angelo Cangelosi. 2019. "Sensorimotor Input as a Language Generalisation Tool: A Neurobotics Model for Generation and Generalisation of Noun-Verb Combinations with Sensorimotor Inputs." *Autonomous Robots* 43 (5): 1271–1290.

21 Knowledge Representation and Reasoning

Michael Beetz

Robots are already making large strides in their abilities, but as the generalizable knowledge representation problem is addressed, the growth of robot capabilities will begin in earnest, and it will likely be explosive. The effects on economic output and human workers are certain to be profound.

—Pratt 2015

21.1 Introduction

One of the most impressive cognitive capabilities of humans is the ability to accomplish their everyday manipulation tasks. In most cases, simple and vague instructions such as “set the table,” “bring me something to drink,” or “clean up” suffice to let us know what to do. The behavior that humans generate in order to perform such manipulation tasks is sophisticated, complex, and tailored to the objects they manipulate, their skill level, the context of the task, and the surrounding scene in which the task is to be performed. Accomplishing these tasks also requires humans to avoid common pitfalls such as breaking objects or spilling fluids.

A main challenge in accomplishing a task such as “set the table” is that it is underdetermined. The request does not spell out which objects to put on the table, the arrangement of the objects, where to find the objects, how they look, how they have to be handled, how they can be efficiently carried, or whether there are social conventions on how to grasp and hold them. Consequently, humans must have the knowledge and the reasoning capacity required to close the gaps between what they are explicitly told and what they are expected to do. This knowledge, including commonsense and intuitive physics knowledge, is shared by most humans, which makes it possible for a person to execute a task to the satisfaction of the person requesting it even if the instructions are vague.

By contrast, imagine how hard it must be to write a robot control program for an autonomous household assistant robot that has to accomplish these tasks in different households, with different objects, for different habits and preferences, and under different circumstances, requiring the program to select the most adequate course of action in so many possible contexts.¹

Many different approaches can be taken to generate robot control programs for tasks such as “set the table,” including robot learning (Peters et al. 2016), task and motion planning

(Lynch and Park 2017; Kavraki and LaValle 2016; Chung, Fu, and Kröger 2016; Villani and Schutter 2016), knowledge-based approaches (Beetz et al. 2012, 2016), and combinations of them.

To make our discussion more concrete, we take a look at the knowledge-based approach to robot programming that is illustrated in figure 21.1. The control program of a robotic agent in the knowledge-based approach consists of a generalized plan and a knowledge base of assertions and asserted reasoning patterns, often called axioms and inference rules.

The generalized plan spells out the logic of the implemented action. For the fetch&place task, this means the robot performs the pickup action at the location where it expects the object to be and places the object at its destination. Lots of complexity is hidden by this simple plan structure. For example, in order to be at a certain location, the robot has to navigate there. And if the robot has to change its position—for example, due to a sudden, more urgent request—it has to interrupt the task and return to it later in order to complete it. Another important aspect hidden in the plan is failure detection, recovery, and continuation. In autonomous robot applications requiring goal-directed object manipulation, more than 80 percent of the programs are concerned with competent failure handling.

A key reason why the robot plan is so compact and elegant is that programmers can state action parameters vaguely. The term “at the location of the object” abstracts away from various pieces of detailed information that a robot needs to perform the task success-

Generalized plan
<p>def-plan fetch&place ((obj): (an object (type thing)), (dest): (a location (type place)))</p> <p>1. at-location (a location (location-of (obj))) perform (an action (type fetching) (object-acted-on (obj)))</p> <p>2. at-location (a location (dest)) perform (an action (type placing) (object-acted-on (obj)) (destination (dest)))</p>
Knowledge base
<ul style="list-style-type: none"> • Cups used for tablesetting have to be clean and unused • People want to use their preferred items • Cups in cupboards are clean • Clean cups are empty • Cups have to be grasped outside • . . .

Figure 21.1

The knowledge-based approach to robot programming includes two main components: a generalized robot plan and a knowledge base of assertions and rules.

fully. For example, in order to pick the object up, the robot has to look at the object with a camera pose that enables it to estimate the pose of the object accurately enough given the inaccuracies of the cameras and occlusions caused by other objects. Then robots often have to reposition themselves to reach the object with the appropriate hand pose, given bulky robot arms. Not specifying these information pieces puts the burden on the robot control programs to infer them automatically.

The programmers also need not specify *how* the object is to be picked up. But consider the scene in figure 21.2 where the object to be picked up is a pot filled with boiling vegetables and water sitting on a hot stove in order to pour the water out. Any robot plan that competently and robustly picks up the pot with the generalized plan has to make the following inferences. It has to infer the motion parameters and constraints for the pickup action, including that the pot has to be picked up with two hands, grasped by its handles, and held horizontally. It has to infer that the handles must be grasped so the robot can tilt the pot around the axis between the handles, that the weight of the pot will change while pouring, and that the lid must be removed before pouring. Finally, the robot has to infer many motion specifics, such as the positions of the robot grippers on the object, the grasp type, and the grasp and lift force as well as the reaching trajectories for the hands.

In order to fill the knowledge gaps, the plan is complemented with the knowledge depicted in the lower part of figure 21.1. This states very general knowledge chunks including facts, rules, and other relationships between objects, tasks, environments, capabilities, and preferences that are asserted to be true. Using this knowledge, the robot can execute an underdetermined action by inferring the appropriate motion parameterization by applying the knowledge in the knowledge base to the given action description in the specific situation's context, as suggested in the example above of picking up the pot from the stove.

Advantages of the knowledge-based approach to robot control over other approaches include the fact that knowledge can be combined by automated reasoning engines in order



Figure 21.2

Easy for humans but difficult for robots: picking up a pot filled with boiling vegetables in order to pour the water out.

to achieve open question-answering capabilities. The abstract format of the knowledge ensures that it can be applied to future situations that are unknown at the time of specification. So, if the robot knows that all open and filled containers have to be held upright, it can use this knowledge for all containers it ever encounters regardless of form and size.

While knowledge-based programming is attractive because of its potential scalability toward open-task domains, it also raises difficult open-research questions. For example, it remains to be seen whether robots can fully leverage knowledge bases in which all knowledge pieces have preconditions that have to be known for the knowledge piece to be applicable. For example, the knowledge that containers have to be held upright is helpful only if the robot can reliably recognize containers. Unfortunately, many components of robot control programs can only provide uncertain information.

There is substantial evidence that accomplishing manipulation tasks requires robotic agents, as well as the human brain, to employ a combination of learning, planning, and other reasoning methods.

21.2 Body Motion Query

Perhaps the most essential reasoning task for a robotic agent manipulating objects is figuring out how to move its body in order to achieve some goal by causing some desired effects and avoiding unwanted side effects. Wolpert (2011), a leading neuroscientist investigating human motion control, argues, “We have a brain for one reason and one reason only, and that is to produce adaptable and complex movements. . . . Movement is the only way you have of affecting the world around you.”

For goal-directed object manipulation, we reason not only about the motions but also about related aspects of actions. These aspects include the relationship between motion parameterization and physical effects, the information preconditions needed to perform actions, and expectations about effects and possible failure modes. If one pours pancake batter onto a pancake maker, the shape of the pancake, whether round or oval or in one piece or more, as well as whether batter will be spilled, depends on the pouring motion. A robotic agent has to acquire knowledge and reason to determine the body motions that enable it to modify its physical surroundings to achieve its goals.

Imagine that a robotic agent is given the task “pour the water out,” which might be stated as a formal expression of the following form:

(perform (an = action (type pouring) (theme (some = substance (type *water*))).

Any robot control program that is to accomplish this underdetermined instruction—including the one depicted in figure 21.2—has to infer *how* the motion is to be generated. As we have argued before, it has to infer the need to grasp the pot by the handles, hold it horizontally, tilt the pot around the axis between the handles, and adjust the force with which to hold the pot according to the changing weight.

Figure 21.3 shows why it is so important to reason about the motions that a robot intends to perform in order to achieve the desired effects of a manipulation task and avoid the unwanted ones. The figure displays examples of unwanted side effects caused by inappropriate motion parameterizations: We see the spatula stuck inside a pancake because the robot did not push the spatula hard enough to slip under it (*upper left*). The robot is not able to lift



Figure 21.3

Action failures and unwanted effects caused by inappropriate motion parameterizations.

the pancake because it has targeted the top of the pancake with the spatula rather than sliding the spatula underneath (*upper right*). The robot has poured too much pancake batter, causing it to spill down the side of the pancake maker (*lower left*). The pancakes are not being properly placed on the plate, causing them to fall off when the plate is lifted (*lower right*).

The composition of elementary movements into the complex movements needed to accomplish actions such as picking and placing and pouring has been investigated in several research areas (Schmidt 1975; Schack et al. 2016), including action science (Prinz, Beisert, and Herwig 2013). Flanagan, Bowman, and Johansson (2006) proposed conceptualizing the action category-specific patterns of movements as motion plans that implement an action as a partially ordered and synchronized set of motion phases. Each motion phase has motion goals, and the transition between motion phases is initiated through perceptually distinctive force-dynamic events (see figure 21.4; Siskind 2001). The motion phases also have knowledge preconditions: in order to execute a reaching motion, I have to know the destination of the reach and the type of grasp to be executed. Thus, to execute a motion plan, the knowledge preconditions of the motion plans have to be inferred.

In order to execute the motion plan for picking up an object, the robot has to infer the body pose with which to start the activity. The robot typically has to be able to see and reach the object. If the object is inside a container, the container often must be opened to reach the object. When starting the reaching motion, the robot must commit to a grasp type, contact points, and a reaching trajectory. These decisions might require the robot to simulate its action and motion plan before carrying it out in the real world. If there are constraints, such as keeping the object upright, or the placement of the object requires the grasp pose to satisfy additional constraints, the robot must foresee the consequences of its parameterization decisions on these future constraints. It also has to decide on the forces it intends to use for grasping and lifting the object.

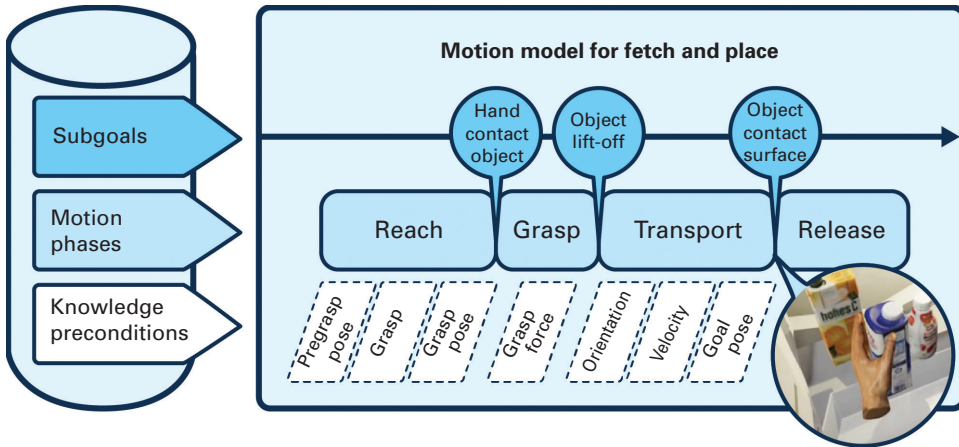


Figure 21.4
Generalized motion plan for a fetch action.

According to this model of implementing actions through movement plans, a promising approach to organizing the computational process for executing underdetermined actions is the following: The robot infers belief about where the object could be found and formulates the instruction with this belief. It then augments the action description with placeholders for the motion parameters and then asks the robot's reasoning system to infer the appropriate parameterizations. This is done by asking the *body motion query*:

how = do I have to move my body
in order to
accomplish the given action description
for the current task
with the objects and in the context
that I see or believe

Answering the body motion query is a very complex and challenging reasoning task. Depending on the context, it might require predicting the physical effects of actions, having commonsense, understanding intuitive physics, knowing social norms, and having experience. In the next section, we will consider how this knowledge can be stated and reasoned about using symbolic knowledge representation.

21.3 Complementary Ways of Structuring Actions

In the previous section, we learned one particular perspective on actions—namely, underdetermined action descriptions and how they can be used to help the robotic agent generate the motions that accomplish goals and avoid unwanted side effects. In this section we look at other perspectives that take complementary views and facilitate other modes of reasoning actions that complement the mechanism introduced in the previous section (Zech et al. 2019).

The first one is to represent and reason about actions by modeling the structure of actions using grammar for understanding and generating natural language. The grammar view

of actions provides a powerful way of dealing with the variations of behaviors and of implementing actions depending on different contexts. The grammatical structure is used for understanding, executing, and learning actions. An example is the grammar proposed by Pastra and Aloimonos (2012) that generates action structures guided by the objects acted on and the tools used.

Another view is to categorize actions and model action categories with respect to the entities that participate in actions and the role they take. For example, in a pouring action you might have a substance that is poured, a container as a source that it is poured from, a destination that it is poured into, and the purpose of the pouring action. In this view we can model action categories as graph structures where the nodes represent the concepts of the entities that participate in the action and links the role that the entities take. Online knowledge services such as FrameNet and VerbNet provide these representations that can be used by robotic agents in order to refine and disambiguate action descriptions (Kipper et al. 2008). Nyga and Beetz (2018), for example, learned joint probability distributions over these graph structures from instructions on websites such as wikiHow that let robots compute the most probable completion given a partial action description as evidence.

Force dynamics (Talmy 1988) is another linguistically motivated approach to represent the structure of actions. It focuses on how entities involved in an action interact with respect to the forces they exert during the action. Force dynamics introduces concepts such as the exertion of, the resistance against, and the blockage of forces, which model the causal structure of an action in a more fine-grained manner.

In addition to linguistically motivated structures of actions, one can also model actions based on body poses and motions. Examples of this line of action modeling are taxonomies of manipulation actions based on hand-object relations (Wörgötter et al. 2013) and the whole-body support taxonomy based on multicontact motions (Borràs et al. 2017). These models have primarily been used for action understanding and imitation learning (Aksoy et al. 2017).

Another approach is the categorization of action categories in terms of the general structures in the sensor data and motion streams generated through the execution of actions. Important models in this dimension are object-action complexes (OACs; Krüger et al. 2011, Wörgötter et al. 2015), which model actions through state prediction functions and probabilistic success measures. The OAC representation is designed to be learned in a bootstrapping fashion and to provide a universal representation for the efficient planning and execution of goal-directed actions.

The structure of actions often becomes more complicated when the environment in which the action takes place is more complex and cluttered and when the changes of the scene are correlated with the success and failure of actions (Yang et al. 2013).

21.4 Symbolic Knowledge Representation and Question Answering

In the early days of artificial intelligence (AI), researchers proposed a powerful class of mechanisms for automating reasoning called *physical symbol systems*. Physical symbol systems are information-processing systems that operate on symbols, combine them into composite symbol structures, and manipulate them to produce new symbol structures (Newell and Simon 1976). They thereby evolve collections of symbol structures by adding, deleting, and modifying them.

One powerful application of physical symbol systems is their use as knowledge representation and reasoning systems: symbol structures can be used as internal representations of knowledge about robots' tasks and actions, and the creation of new symbol structures can be used to draw conclusions from the knowledge. Using physical symbol systems, programmers can equip robots with symbol structures representing the tasks that a robot is to accomplish, the actions it can execute, the environment it is acting in, the objects that it manipulates, and their states. Physical symbol systems can then implement intelligent reasoning, decision-making, and planning as mechanical symbol manipulations.

One of the most prominent categories of physical symbol systems is logic (Hayes 1977).² A logic consists of three components: its syntax, its semantics, and its calculus. The syntax of the logic defines what can be expressed; it is the set of symbolic expressions that constitute the language of the logic. The semantics assigns truth values to expressions—that is, it defines whether expressions are true or false. Finally, the calculus defines the rules for creating new symbol structures out of existing ones. Thus, logics are physical symbol systems in which the semantics defines whether a given symbol structure is true or false.

In order to use logics for implementing reasoning in computer systems, and in particular for enabling robots to decide on their courses of action, researchers aim to design logics with which they can express relevant problems and the knowledge that is necessary to solve them. In addition, they define a semantics for the new logics in which the truth values are defined on the basis of the truth values of the constituent structures. Finally, they aim at defining a calculus in which a symbol structure can be derived from a set of symbol structures if and only if the derived symbol structure is true if the original symbol structures are true. In this case the calculus is called correct and complete with respect to the semantics of the logic.

Logics with correct and complete calculi are potentially very powerful tools for solving problems with computers. They allow computers to solve problems without requiring the computer to understand the domain of problem-solving. This is possible because solving a problem p can be implemented as answering the question of whether there exists a solution for p . To determine the answer, it suffices to determine whether the statement “there exists a solution for p ” is true. In a logic with a complete and correct calculus, this is exactly the case if we can derive the symbolic expression that represents “there exists a solution for p .”

This gives us a method for solving problems that can be automated in a straightforward manner: Given a set of symbolic expressions that are asserted to be true, generate all symbolic structures that can be derived using the rules of the calculus. If for the problem p that we want to solve the symbol structure representing “there exists a solution for p ” is in the resulting collection of symbol structures, then we know a solution exists.

Some logic calculi have technical properties that make them particularly attractive for problem-solving. The first such property is that the tree of derivations that result in the symbol structure representing the solution constitutes a rigorous proof of the existence of a solution. Some calculi also provide the proof of existence by generating an example solution, which is what we want in the first place. The proof can also be used as an explanation of why the generated example is a solution, which is an asset for constructing explainable AI systems.

From this perspective, predicate logic, together with some of its calculi, is a particularly powerful and adequate logic (Kowalski 1979). Predicate logic is a logic capable of express-

ing factual knowledge of our natural language. Together with the resolution, calculus can be used as a problem-solver. The programming language Prolog, which is a subset of predicate logic and provides a limited implementation of negation, is a pragmatic alternative for programming problem solvers, which suffices for most of our purposes.

Taking the logic approach gives us a potentially very elegant and powerful way to program robotic agents, which is illustrated in figure 21.5. In the lower part, we see a robot accomplishing everyday manipulation tasks such as setting and cleaning the table. Researchers propose to distill the knowledge that the robot needs to accomplish its tasks as a collection of symbolic expressions asserted to be true, which is called the axiomatization of the problem-solving domain. Each symbolic expression, called an axiom, corresponds to a piece of human knowledge, and this correspondence is implied by the semantics of the logic. Researchers then aim at inventing axiomatizations in a correct and complete calculus of a problem-solving domain that are strong enough to solve all problems in this domain. This means that whenever the robot needs to achieve a goal g starting from the current state s , this task can be transformed into a logical-reasoning problem that can be solved mechanically. In order to do this, the environment, the current state, models of the actions, and other kinds of knowledge have to be asserted as axioms. The question “Does there exist a state of the environment that satisfies the goal and can be reached through a sequence of actions?” must be asked. Then the calculus proves the existence by searching for a symbolic expression that represents such a state. The way this state can be achieved—that is, the sequence of actions that has to be executed—is contained in the existence proof. Now the real goal

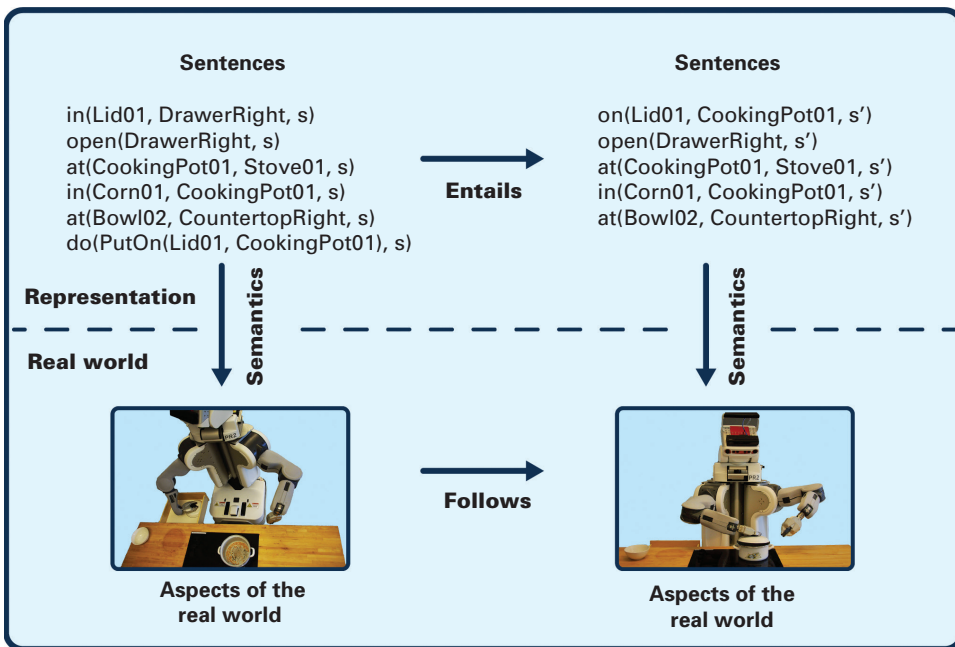


Figure 21.5
Deciding on a course of action using logic-based inference engines.

state and the actions the robot is to execute can be computed as the meaning of the respective logic expressions.³

Now suppose that we can axiomatize the actions that a robot can perform, the conditions under which the actions are executable, and their physical effects such that each executable action sequence and the state resulting from the sequence can be inferred from the axioms. This set of axioms is a valuable knowledge source for robotic agents that are to perform open tasks in open domains. Using the axioms, the robotic agents can reprogram themselves to accomplish new tasks. The action sequences they return are proven to achieve the given goals, and they can generate an action sequence for every task they are capable of doing.

A large community of researchers has followed this research direction. McDermott proposed an axiomatization of problem-solving in a first-order time interval logic by providing a powerful set of axioms talking about plans and their execution and the physical effects they cause. Hayes (1968, 1979) proposed a comprehensive research enterprise aiming at formalizing the commonsense and naive physics knowledge and reasoning needed to solve a broad range of everyday tasks. This research direction has also been put forward in textbooks by (Davis 1990) and (Mueller 2006).

Many researchers have proposed component axiomatizations for specific categories of reasoning problems, including reasoning about actions (Reiter 2001); qualitative reasoning (Davis 2017; Davis and Marcus 2015); spatial and temporal reasoning (Allen 1984); constraint and resource reasoning; rational agency by formalizing the relations between the beliefs, desires, and intentions of agents (Georgeff et al. 1998; Rao and Georgeff 1992); and multiagent activity (Hoek and Wooldridge 2012; Wooldridge 2009).

The logic approach to problem-solving has also raised some questions regarding its feasibility. One of these questions is whether we can find general calculi that solve all relevant problems. This was originally brought up in the context of reasoning about actions. When you try to predict what will happen, you typically want your inference system to have the bias that changes in the world only occur if they are forced to. In other words, the world has the tendency to stay as it is, and change tends to occur as late as possible (the law of inertia). If you are reasoning backward in order to explain why a change occurred, it does not make sense to assume that the change occurred immediately before noticing it. This seems to suggest that different inference processes are needed depending on the question you ask—whether you reason forward or backward in time (Hanks and McDermott 1987; McDermott 1987). Another essential problem is that symbolic representations represent objects and states in the world. So imagine that the symbol structure *cup-23* stands for my cup. If a robotic agent looks at a table with two identical cups sitting next to each other, the robotic agent might not have the perceptual ability to distinguish my cup from the other one and therefore is not able to execute the action *pick-up (my cup)*. This and related problems are often referred to as the *symbol-grounding problem* (Harnad 1990), and some variations of the problem are addressed through reasoning about the knowledge preconditions of plans (Moore 1984; Morgenstern 1987).

Occasionally, reasoning challenges have been proposed that require a combination of different reasoning capabilities, including spatial and temporal reasoning and reasoning about action and change. One of these challenges is the egg-cracking problem (Miller and Morgenstern 1997; Morgenstern 2001). It is based on a sequence of actions that leads to

an egg being cracked and the egg yolk to be separated from the egg white and dropped into a bowl. The challenge for logical reasoning is to answer an open set of “what if” questions, including what happens if the egg is hit on the table very smoothly or very forcefully, if the egg is from an ostrich, if the bowl is placed upside down, and so on. Can we formalize a compact axiom set that entails all the answers to these “what if” questions? As it turns out, the axiom sets become huge quickly, and the appropriate level of abstraction depends on the question to be answered.

Another problem is the effectiveness and efficiency of the reasoning processes. If axiomatizations are very comprehensive and general, the axioms can be used in many ways to generate new symbolic structures, and the search space for a proof can be highly exponential and exceed the available resources. Therefore, more efficient representations and algorithms have been proposed and investigated in order to infer action plans for robotic agents. Here the representations of actions for planning (Fikes and Nilsson 1971; Ghallab et al. 1998; Fox and Long 2011) and special-purpose planning algorithms are particularly important for robotics applications (Ghallab, Nau, and Traverso 2004, 2016).

Unfortunately, the capability of inferring provably correct plans does not mean that the plans will work when executed. This is because the axioms formalize idealized models of the world, robot capabilities, and actions and their effects. One reason why we cannot equip robotic agents with faithful logic models of their perception and action capabilities is that perception, action, and, consequently, robot beliefs about the world are incomplete and uncertain. One prominent way to competently reason with uncertainty is to use probabilistic representation and reasoning methods (Thrun, Burgard, and Fox 2005).

21.5 Ontologies and Encyclopedic Knowledge Bases

Robots need comprehensive knowledge about their tasks, their bodies and capabilities, the objects they are to manipulate, and their environments. An outdoor drone might want to use a web service such as OpenStreetMap as an information resource for landmarks, street maps, or building functions. A robot that is loading and unloading machines in a factory might need structural, functional, and process knowledge about the machine to act more competently. A key question for deploying cognitive robots for a variety of tasks in a variety of domains is how the necessary knowledge can be provided in a structured and organized manner (Noy and McGuinness 2013).

Key mechanisms in knowledge representation and reasoning that have been developed for these purposes are encyclopedic knowledge bases, particularly in the form of ontologies (Baader et al. 2007). Ontologies specify the concepts—categories of entities—needed for answering questions about a problem domain in a machine-understandable manner such that symbolic reasoning methods can use them. For example, in an ontology you can specify the concept “refrigerator” as specializations of the concepts “electrical device” and “container.” Because all instances of the concept “electrical device” were previously defined to have a property “state,” which can take the values “switched on” or “switched off,” all refrigerators inherit this property because refrigerators were defined as specializations of electrical devices. Because refrigerators are also containers, we know from the concept description of “container” that they have a capacity and can be opened. We might want to assert additional properties for the concept “refrigerator”—namely, that the primary function

of a refrigerator is to store perishable food. Here perishable food is again a concept defined to be a specialization of the concepts “perishable items” and “food.”

Now, if the perception system of a robot categorizes an object as a refrigerator, it can assert that the detected object is an instance of the concept “refrigerator.” By making this assertion, the robot automatically infers that the detected object satisfies all the knowledge it has about a refrigerator. In particular, if the robot searches for milk and knows that milk is a perishable food, it can automatically infer that it might be able to find the milk inside the refrigerator because that is a storage place for perishable food.

More generally, the key idea of an ontology is to name, define, and formally represent concepts in terms of more primitive concepts, their properties, and their relations to other concepts. The collection of defined concepts is the vocabulary that can be used to represent and reason about an application domain.

Of course, such general knowledge is applicable to different tasks and environments. This fact has motivated research with the goal of developing a comprehensive and common ontology that can serve many different applications. Perhaps the best-known ontology that has been developed for this purpose is the Cyc ontology (Lenat 1995). The Cyc knowledge base was developed more than thirty-five years ago. Cyc contains an ontology of about 1.5 million general concepts and more than 25 million general rules and assertions involving these concepts and representing how the world works. The Cyc knowledge base includes commonsense knowledge and knowledge that is typically implicit.

Ontologies, developed in the knowledge representation field, are mostly developed for question answering and problem-solving applications. Therefore, they are typically too abstract for robot control and have to be extended through the addition of domain-specific subontologies to cover robot agency (Olivares-Alarcos et al. 2019).

Ontologies are also used to make information available in the internet machine-understandable. An area investigating this research direction is called the semantic web technology (Hendler 2001; Heflin and Hendler 2001). Its basic idea is to represent the information contained on a web page as logical fact. Then the formatting of the web page is automatically generated through rules specified by the web programmers. Names of the predicates used in the logical representation and the categories of objects used as terms are defined in ontologies that the web pages point to. By reading the logical facts of the web page and the corresponding ontology, computer programs can automatically reason about the information contents of web pages.

Potentially, the semantic web is a powerful enabler for cognitive robots. Imagine that a retailer has a web store implemented as a semantic website. A cognitive service robot acting in the physical retail store could use the content of the semantic web store as part of its domain model.

Unfortunately, only a small part of the internet is encoded using semantic web technology. Therefore, researchers have started to automatically construct knowledge bases through statistical learning based on huge amounts of web data. This way a computer program can learn that Stephen Curry is a basketball player, plays for the Golden State Warriors, and plays in the position of a guard. Knowledge bases built this way are called lightweight knowledge bases. Perhaps the most important knowledge base built this way is the Google knowledge graph, which reportedly included more than seventy billion facts by the year 2016.

As knowledge bases acquired through statistical learning from large databases are built on correlations rather than facts, they may contain inconsistencies and faulty knowledge pieces. Because logic-based reasoning engines cannot deal with inconsistent knowledge, other forms of reasoning are needed. To allow for possible inconsistencies in knowledge bases, scalable inference systems employ an ensemble of expert reasoning methods that do heuristic reasoning. Results from the individual reasoning mechanisms are then treated as hypotheses for possible answers and solutions, which have to be tested and rated according to their plausibility. The IBM Watson system that won the famous quiz show *Jeopardy* against human champions is a successful example of this technology (Ferrucci et al. 2010). Given a quiz question, the system generates hundreds of possible answers and tests and ranks them according to plausibility within three seconds.

For cognitive robots, hybrid knowledge representation and reasoning seems the most promising path to go. This means using correct and well-designed knowledge representation and reasoning where possible and employing the huge data and information resources available in many modern information services. For the successful agency of robotic agents, it is of key importance that existing knowledge sources are combined with the robot's own experiences to make the knowledge actionable and tailor it to the robot's needs.

21.6 Knowledge Representation and Reasoning Systems for Cognitive Robots

Several knowledge representation and reasoning systems have been specifically designed for autonomous robot control, including KnowRob (Tenorth and Beetz 2013; Tenorth and Beetz 2015; Beetz et al. 2018), ORO (Lemaignan et al. 2010), and ROSETTA (Topp et al. 2018). In this section we discuss the specifics of robot knowledge representation and reasoning in the context of the KnowRob system, which is the most comprehensive and widely used robot knowledge representation and reasoning system (Olivares-Alarcos et al. 2019).

The KnowRob KR&R system is open source, with documentation, installation guides, and tutorials at www.knowrob.org. KnowRob can also be used through the web-based knowledge service www.open-ease.org.

The software architecture of the KnowRob KR&R system is depicted in figure 21.6. The core of the system is an ontology that makes all concepts needed for question answering and reasoning explicit and machine-understandable. The ontology is surrounded by the hybrid reasoning kernel, which employs several mixed symbolic-subsymbolic representation structures. This kernel is surrounded by the logic-based language shell, which provides a unified symbolic interface to the different representations and reasoning methods of the hybrid reasoning kernel. The outer layer consists of the interfaces to the different applications of KR&R for robot control, including question answering, perception, the recording of experiences, and the learning from knowledge and experience.

Compared to KR&R systems for other application domains, those for autonomous robot control employ ontologies that are particularly strong and expressive with respect to the representation of objects and actions. Object representations typically include physical and

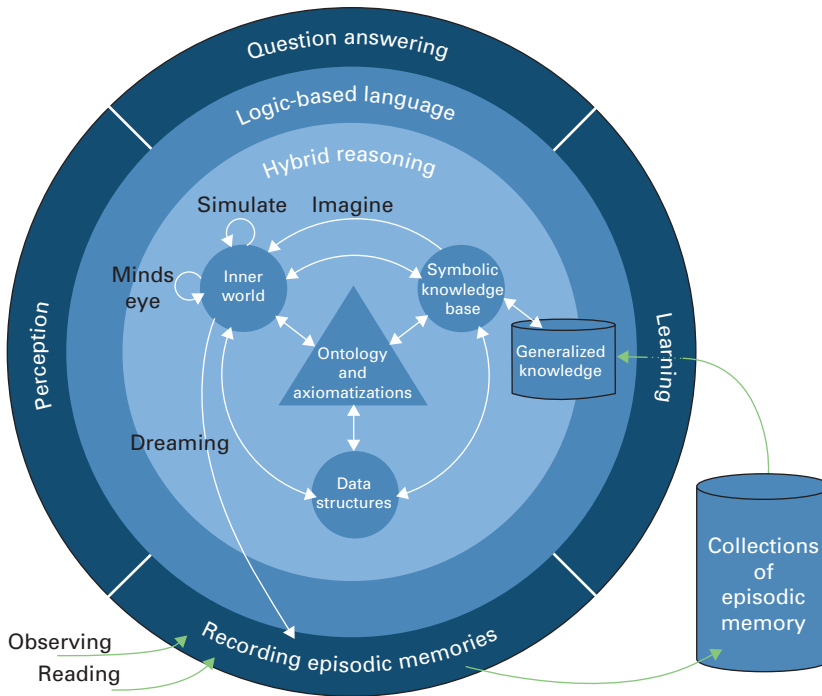


Figure 21.6
Software architecture of the KnowRob KR&R system.

geometric properties, their parts, articulation models, functional information, and visual appearance. The representation of actions often includes the behavior they generate, the physical effects they cause, and the intentions they serve.

Another particularity is that the represented domain is accessible to the KR&R system because the representation and the represented control system reside in the same computer system. Therefore, robot programmers can develop a robot capable of using its control system and perception system as knowledge sources. Often, information that is needed for abstract reasoning is already available in some form in the robot’s internal data structures, such as the robot’s pose estimate, or can be acquired from its components, such as the perception system. To reuse this information, the robot can “listen” to the control program, record the dynamic data structures, and use the data as a dynamic and virtual knowledge base (Mösenlechner, Demmel, and Beetz 2010). The knowledge-processing system thus reuses and abstracts data structures used by the control program for the purpose of reasoning. Since the knowledge is generated on demand and just in time, the abstract representations are solidly grounded in these data structures.

A second representation and reasoning component is the inner-world knowledge base (Ziemke 2001). It is a detailed and photorealistic reconstruction of the robot’s environment in a game engine with physics simulation and vision capabilities and adds powerful reasoning methods to the knowledge-processing framework. First, the robot can geometrically reason about a scene by virtually looking at it using the vision capability provided by the game engine (Qiu and Yuille 2016) and predict the effects of actions through semantic

annotations of force dynamic events monitored in its physics simulation. As Winston (2012) would phrase it, it allows the robot to reason with its eyes and hands. All physical entities in the game engine are also entities in the symbolic knowledge base, which means the game engine state is correctly, accurately, and completely represented by the inner-world knowledge base.

A third component is the symbolic knowledge base that contains common-sense, intuitive physics knowledge as well as domain knowledge. In many cases the domain knowledge bases can be constructed automatically from semistructured web pages such as Google Maps, OpenStreetMap, web stores, DBpedia and the likes.

As the representations and mechanisms used in the hybrid reasoning shell were originally created for the action execution shell and are redundant, the solutions hypothesized by the individual methods may have to be checked and ranked with respect to their plausibility as, for example, was done in the Watson system (Ferrucci et al., n.d.).

The subsequent interface layer casts the hybrid reasoning shell as a first-order logic knowledge base that is largely constructed on demand from data structures of the control program and computed through robotics algorithms.

21.7 Neurosymbolic Learning and Reasoning

In recent years the machine-learning approach to perception, action, and intelligent problem-solving has gained a lot of momentum (Hassabis et al. 2017). In particular, deep learning and deep reinforcement learning (Silver et al. 2016; Berner et al. 2019) have achieved impressive successes in specific tasks of autonomous agency. Where symbolic knowledge representation has its strengths in generalization, modular and compositional structure, and potential for introspective capabilities, artificial neural representations have their strengths in learning representations that are well correlated and in learning complex action selection, question answering, and problem-solving tasks in an end-to-end fashion (Levine et al. 2016, 2018; Sünderhauf et al. 2018). Technically speaking, action selection and execution are computational tasks that map the continuous sensor data streams into continuous motion actuation functions and therefore are in the applicability domain of deep network technology. However, limitations have been identified in the robust handling of problem instances that are not covered in the experience data, in explaining and diagnosing the generated behavior in order to quickly adapt to unpredicted circumstances, and in efficient learning from little experience (Marcus and Davis 2019).

Inspired by these successes and considerations, some researchers have proposed methods to combine the strengths of both approaches and extend neural representations with operations that replicate some of the functionality of symbolic representations or combine artificial neural learning mechanisms with symbolic reasoning. The characteristics of these approaches show substantial promise for representing the structures of actions and reasoning about them. Examples of such learning and reasoning approaches include hyperdimensional (Kanerva 2009, 2018) and neurosymbolic computing (Garcez et al. 2019; Besold et al. 2017).

Hyperdimensional computing (Neubert et al. 2019) performed by vector symbolic architectures represents symbols through high-dimensional vectors (typically thousands of dimensions) and exploits the redundancy of the encoding to achieve robustness to noise

and uncertainty. In addition, it employs operators to perform symbolic computations with high-dimensional vectors (Gayler 2004; Blouw et al. 2016; Levy and Gayler 2008). These operators enable the encoding of prior knowledge, the generalization of concepts from similar symbols, the composition of complex expression, and thereby also the advantage of learning from fewer examples. Thus, the intuition is to add some representation and reasoning capabilities to high-dimensional vector spaces (Eliasmith et al. 2012). Examples in which hyperdimensional computing is applied to intelligent robot agency include active perception, place recognition, and the learning and recalling of reactive behavior.

Another category of approaches is neural-symbolic computation (Garcez et al. 2019). It aims at integrating robust vector-based learning and symbolic reasoning by implementing new powerful alternatives for knowledge representation, learning, and reasoning based on neural computation.

21.8 Conclusion

Knowledge representation and reasoning will be a key information-processing capability for cognitive robots that are to accomplish vaguely specified tasks in open environments. Knowledge processing can complement machine-learning decision-making and control mechanisms because reasoning steps are based on rules that can be asserted to be valid. Additionally, the use of knowledge representation and reasoning substitutes the black box reasoning of machine-learning methods with justifiable inference chains that make the reasoning transparent and enable cognitive robots to reason about their decision-making. It is essential that the methods not only work in the abstract but apply to the sensory and motion level to achieve the full potential of the representation and reasoning methods. Leveraging modern information-processing techniques—including realistic simulation and rendering techniques, neurosymbolic and hyperdimensional computing, and big data and data-intensive machine-learning methods—provides promising opportunities to do this.

Additional Reading and Resources

- A seminal, extensive volume on commonsense knowledge representation: Davis, Ernest. 1990. *Representations of Commonsense Knowledge*. The Morgan Kaufmann Series in Representation and Reasoning. Burlington, MA: Morgan Kaufmann. ISBN 978-1-55860-033-1.
- A comprehensive presentation of the KnowRob robot knowledge representation and reasoning architecture: Tenorth, Moritz, and Michael Beetz. 2013. “KnowRob: A Knowledge Processing Infrastructure for Cognition-Enabled Robots.” *International Journal of Robotics Research* 320 (5): 566–590. <http://ijr.sagepub.com/content/32/5/566.short>.
- A systematic taxonomy for categorizing action representations in robotics along various dimensions, with a meticulous literature survey on action representations in robotics: Zech, Philipp, Erwan Renaudo, Simon Haller, Xiang Zhang, and Justus H. Piater. 2019. “Action Representations in Robotics: A Taxonomy and Systematic Classification.” *International Journal of Robotics Research* 380 (5). <https://doi.org/10.1177/0278364919835020>.

- The knowledge representation and reasoning framework KnowRob is accessible, including open-source software, documentation, installation guides, and tutorials, at www.knowrob.org. KnowRob can also be used through the web-based knowledge service openEASE: www.open-ease.org. Examples of the application of reasoning to the plan-based control of robotic agents were realized through CRAM (Cognitive Robot Abstract Machine), which is accessible through the website www.cram-system.org.

Notes

1. Kemp et al. (2007) and Mustafa Ersen et al. (2017) provide comprehensive review articles about challenges and approaches to autonomous robot manipulation in human environments.
2. Excellent textbooks on the logic-based approach to building intelligent systems include Genesereth and Nilsson (1987); Reiter (2001); Davis (1990).
3. As stated in chapter 1, this approach constitutes one of the roots of cognitive robotics in which Levesque, Reiter, De Giacomo, Lakemeyer, and colleagues propose to model high-level robotic control using explicit knowledge and reasoning in order to decide which actions to execute (Levesque and Lakemeyer 2008). This chapter adopts this view but does not limit the representation of actions and reasoning to the high level of abstraction. Rather, we extend the view to reasoning about how actions should be executed.

References

- Aksoy, Eren Erdal, Adil Orhan, and Florentin Wörgötter. 2017. "Semantic Decomposition and Recognition of Long and Complex Manipulation Action Sequences." *International Journal of Computer Vision* 122 (1): 84–115.
- Allen, James F. 1984. "Towards a General Theory of Action and Time." *Artificial Intelligence* 23 (2): 123–154. doi:10.1016/0004-3702(84)90008-0.
- Baader, Franz, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider. 2007. *The Description Logic Handbook*. Cambridge: Cambridge University Press.
- Beetz, Michael, Daniel Beßler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoglu, and Georg Bartels. 2018. "KnowRob 2.0—a 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents." In *International Conference on Robotics and Automation*, 512–519. Brisbane, Australia : IEEE.
- Beetz, Michael, Raja Chatila, Joachim Hertzberg, and Federico Pecora. 2016. "AI Reasoning Methods for Robotics." In Siciliano and Khatib 2016, 329–356. doi:10.1007/978-3-319-32552-1_14.
- Beetz, Michael, Dominik Jain, Lorenz Mösenlechner, Moritz Tenorth, Lars Kunze, Nico Blodow, and Dejan Pangercic. 2012. "Cognition-Enabled Autonomous Robot Control for the Realization of Home Chore Task Intelligence." *Proceedings of the IEEE* 100 (8): 2454–2471.
- Berner, Christopher, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, et al. 2019. "Dota 2 with Large Scale Deep Reinforcement Learning." Arxiv preprint: 1912.06680.
- Besold, Tarek R., Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, et al. 2017. "Neural-Symbolic Learning and Reasoning: A Survey and Interpretation." ArXiv preprint: 1711.03902.
- Blouw, Peter, Eugene Solodkin, Paul Thagard, and Chris Eliasmith. 2016. "Concepts as Semantic Pointers: A Framework and Computational Model." *Cognitive Science* 40 (5): 1128–1162.
- Borràs, Júlia, Christian Mandery, and Tamim Asfour. 2017. "A Whole-Body Support Pose Taxonomy for Multi-contact Humanoid Robot Motions." *Science Robotics* 2 (13).
- Chung, Wan Kyun, Li-Chen Fu, and Torsten Kröger. 2016. "Motion Control." In Siciliano and Khatib 2016, 163–194. doi:10.1007/978-3-319-32552-1_8.
- Davis, Ernest. 1990. *Representations of Commonsense Knowledge*. The Morgan Kaufmann Series in Representation and Reasoning. Burlington, MA: Morgan Kaufmann.
- Davis, Ernest. 2017. "Logical Formalizations of Commonsense Reasoning: A Survey." *Journal of Artificial Intelligence Research* 59:651–723. doi:10.1613/jair.5339.
- Davis, Ernest, and Gary Marcus. 2015. "Commonsense Reasoning and Commonsense Knowledge in Artificial Intelligence." *Communications of the ACM* 58 (9): 92–103. doi:10.1145/2701413.

- Eliasmith, Chris, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis Dewolf, Yichuan Tang, and Daniel Rasmussen. 2012. "A Large-Scale Model of the Functioning Brain." *Science* 338 (6111): 1202–1205.
- Ersen, Mustafa, Erhan Oztop, and Sanem Sariel. 2017. "Cognition-Enabled Robot Manipulation in Human Environments: Requirements, Recent Work, and Open Problems." *IEEE Robotics and Automation Magazine* 24 (3): 108–122.
- Ferrucci, David, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A. Kalyanpur, Adam Lally, et al. 2010. "Building Watson: An Overview of the Deepqa Project." *AI Magazine* 31 (3): 59–79. <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2303>.
- Fikes, Richard, and Nils J. Nilsson. 1971. "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving." *Artificial Intelligence* 2 (3/4): 189–208. doi:10.1016/0004-3702(71)90010-5.
- Flanagan, J. Randall, Miles C. Bowman, and Roland S. Johansson. 2006. "Control Strategies in Object Manipulation Tasks." *Current Opinion in Neurobiology* 16 (6): 650–659.
- Fox, Maria, and Derek Long. 2011. "PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains." CoRR Abs/1106.4561. ArXiv preprint: <http://arxiv.org/abs/1106.4561>.
- Garcez, Artur d'Avila, Marco Gori, Luis C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. 2019. "Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning." ArXiv preprint: 1905.06088.
- Gayler, Ross W. 2004. "Vector Symbolic Architectures Answer Jackendoff's Challenges for Cognitive Neuroscience." ArXiv preprint: Cs/0412059.
- Genesereth, Michael, and Nils Nilsson. 1987. *Logical Foundations of Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann.
- Georgeff, Michael, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. 1998. "The Belief-Desire-Intention Model of Agency." In *International Workshop on Agent Theories, Architectures, and Languages*, 1–10. Berlin: Springer.
- Ghallab, Malik, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. 1998. "PDDL—The Planning Domain Definition Language." <https://www.csee.umbc.edu/courses/671/fall12/hw/hw6/pddl1.2.pdf>.
- Ghallab, Malik, Dana S. Nau, and Paolo Traverso. 2004. *Automated Planning—Theory and Practice*. San Diego: Elsevier.
- Ghallab, Malik, Dana S. Nau, and Paolo Traverso. 2016. *Automated Planning and Acting*. Cambridge: Cambridge University Press. <http://www.cambridge.org/de/academic/subjects/computer-science/artificial-intelligence-and-natural-language-processing/automated-planning-and-acting?format=hb>.
- Hanks, Steve, and Drew V. McDermott. 1987. "Nonmonotonic Logic and Temporal Projection." *Artificial Intelligence* 33 (3): 379–412. doi:10.1016/0004-3702(87)90043-9.
- Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D* 42:335–346.
- Hassabis, Demis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. 2017. "Neuroscience-Inspired Artificial Intelligence." *Neuron* 95 (2): 245–258.
- Hayes, Patrick J. 1968. "The Second Naive Physics Manifesto." In *Formal Theories of the Commonsense World*, edited by J. R. Hobbs and R. C. Moore, 1–36. Norwood, NJ: Ablex.
- Hayes, Patrick J. 1977. "In Defense of Logic." In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* 1:559–565. <https://www.ijcai.org/Proceedings/77-1/Papers/099.pdf>.
- Hayes, Patrick J. 1979. "The Naive Physics Manifesto." In *Expert Systems in the Micro Electronic Age*, edited by D. Michie, 242–270. Edinburgh: Edinburgh University Press.
- Heflin, Jeff, and James A. Hendler. 2001. "A Portrait of the Semantic Web in Action." *IEEE Intelligent Systems* 16 (2): 54–59. doi:10.1109/5254.920600.
- Hendler, James A. 2001. "Agents and the Semantic Web." *IEEE Intelligent Systems* 16 (2): 30–37. doi:10.1109/5254.920597.
- Hoek, Wiebe Van der, and Michael J. Wooldridge. 2012. "Logics for Multiagent Systems." *AI Magazine* 33 (3): 92–105. <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2427>.
- Kanerva, Pentti. 2009. "Hyperdimensional Computing: An Introduction to Computing in Distributed Representation with High-Dimensional Random Vectors." *Cognitive Computation* 1 (2): 139–159.
- Kanerva, Pentti. 2018. "Computing with High-Dimensional Vectors." *IEEE Design and Test* 36 (3): 7–14.
- Kavraki, Lydia E., and Steven M. Lavalle. 2016. "Motion Planning." In *Siciliano and Khatib 2016*, 139–162. doi:10.1007/978-3-319-32552-1_7.
- Kemp, Charles C., Aaron Edsinger, and Eduardo Torres-Jara. 2007. "Challenges for Robot Manipulation in Human Environments." *IEEE Robotics and Automation Magazine* 14 (1): 20.

- Kipper, Karin, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. "A Large-Scale Classification of English Verbs." *Language Resources and Evaluation* 42 (1): 21–40.
- Kowalski, Robert A. 1979. "Algorithm = Logic + Control." *Communications of the ACM* 22 (7): 424–436. <http://dblp.uni-trier.de/db/journals/cacm/cacm22.html#kowalski79>.
- Krüger, Norbert, Christopher Geib, Justus Piater, Ronald Petrick, Mark Steedman, Florentin Wörgötter, Aleš Ude, et al. 2011. "Object-Action Complexes: Grounded Abstractions of Sensory-Motor Processes." *Robotics and Autonomous Systems* 59 (10): 740–757.
- Lemaignan, Séverin, Raquel Ros, Lorenz Mösenlechner, Rachid Alami, and Michael Beetz. 2010. "ORO, A Knowledge Management Module for Cognitive Architectures in Robotics." In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3548–3553. New York: IEEE.
- Lenat, Douglas B. 1995. "CYC: A Large-Scale Investment in Knowledge Infrastructure." *Communications of the ACM* 38 (11): 33–38.
- Levesque, Hector, and Gerhard Lakemeyer. 2008. "Cognitive Robotics." *Foundations of Artificial Intelligence* 3:869–886.
- Levine, Sergey, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. 2016. "End-to-End Training of Deep Visuomotor Policies." *Journal of Machine Learning Research* 17 (1): 1334–1373.
- Levine, Sergey, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. 2018. "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection." *International Journal of Robotics Research* 37 (4–5): 421–436.
- Levy, Simon D., and Ross Gayler. 2008. "Vector Symbolic Architectures: A New Building Material for Artificial General Intelligence." In *Proceedings of the 2008 Conference on Artificial General Intelligence*, 414–418. Amsterdam: IOS Press.
- Lynch, Kevin M., and Frank C. Park. 2017. *Modern Robotics: Mechanics, Planning, and Control*. 1st ed. New York: Cambridge University Press.
- Marcus, Gary, and Ernest Davis. 2019. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Pantheon.
- McDermott, D. 1987. "A Critique of Pure Reason." *Computational Intelligence* 3 (3): 151–160.
- Miller, Rob, and Leora Morgenstern. 1997. "Common Sense Problem Page." Stanford University. <http://www-formal.stanford.edu/leora/commonsense/>.
- Moore, Robert C. 1984. "A Formal Theory of Knowledge and Action." In *Formal Theories of the Commonsense World*, edited by J. R. Hobbs and R. C. Moore, 269–317. Norwood, NJ: Ablex.
- Morgenstern, Leora. 1987. "Knowledge Preconditions for Actions and Plans." In *Proceedings of the 10th International Joint Conference on Artificial Intelligence* 2:867–874.
- Morgenstern, Leora. 2001. "Mid-Sized Axiomatizations of Commonsense Problems: A Case Study in Egg Cracking." *Studia Logica* 67 (3): 333–384. doi:10.1023/a:1010512415344.
- Mösenlechner, Lorenz, Nikolaus Demmel, and Michael Beetz. 2010. "Becoming Action-Aware through Reasoning about Logged Plan Execution Traces." In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2231–2236. New York: IEEE.
- Mueller, Erik T. 2006. *Commonsense Reasoning*. Burlington, MA: Morgan Kaufmann.
- Neubert, Peer, Stefan Schubert, and Peter Protzel. 2019. "An Introduction to Hyperdimensional Computing for Robotics." *KI-Künstliche Intelligenz* 33 (4): 319–330.
- Newell, Allen, and Herbert A. Simon. 1976. "Computer Science as Empirical Inquiry: Symbols and Search." *Communications of the ACM* 19 (3): 113–126. doi:10.1145/360018.360022.
- Noy, Natasha, and Deborah McGuinness, eds. 2013. "Final Report on the 2013 NSF Workshop on Research Challenges and Opportunities in Knowledge Representation." National Science Foundation Workshop Report. http://krnsfworkshop.cs.illinois.edu/final-workshop-report/KRChallengesAndOpportunities_FinalReport.pdf.
- Nyga, Daniel, and Michael Beetz. 2018. "Cloud-Based Probabilistic Knowledge Services for Instruction Interpretation." In *Robotics Research*, 649–664. Cham, Switzerland: Springer.
- Olivares-Alarcos, Alberto, Daniel Beßler, Alaa Khamis, Paulo Goncalves, Maki Habib, Julita Bermejo, Marcos Barreto, et al. 2019. "A Review and Comparison of Ontology-Based Approaches to Robot Autonomy." *The Knowledge Engineering Review*. Vol. 34. Cambridge: Cambridge University Press. doi:10.1017/s0269888919000237.
- Pastra, Katerina, and Yiannis Aloimonos. 2012. "The Minimalist Grammar of Action." *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1585): 103–117.
- Peters, Jan, Daniel D. Lee, Jens Kober, Duy Nguyen-Tuong, J. Andrew Bagnell, and Stefan Schaal. 2016. "Robot Learning." In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 357–398. Springer Handbooks. Berlin: Springer. doi:10.1007/978-3-319-32552-1_15.

- Pratt, Gill. 2015. "Is a Cambrian Explosion Coming for Robotics?" *Journal of Economic Perspectives* 29:51–60.
- Prinz, Wolfgang, Miriam Beisert, and Arvid Herwig, eds. 2013. *Action Science: Foundations of an Emerging Discipline*. Cambridge, MA: MIT Press.
- Qiu, Weichao, and Alan Yuille. 2016. "Unrealcv: Connecting Computer Vision to Unreal Engine." ArXiv preprint: 1609.01326.
- Rao, Amand, and Michale P. Georgeff. 1992. "An Abstract Architecture for Rational Agents." In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (Kr'92)*, edited by B. Nebel, C. Rich, and W. Swartout, 439–449. San Mateo, CA: Morgan Kaufmann.
- Reiter, Raymond. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. Illustrated ed. Cambridge, MA: MIT Press.
- Schack, Thomas, Christoph Schütz, André Frank Krause, and Christian Seegelke. 2016. "Representation and Anticipation in Motor Action." In *Anticipation across Disciplines*, edited by Mihai Nadin, 203–215. Berlin: Springer. doi:10.1007/978-3-319-22599-9_13.
- Schmidt, Richard A. 1975. "A Schema Theory of Discrete Motor Skill-Learning." *Psychological Review* 82 (4): 225–260.
- Siciliano, Bruce, and Oussama Khatib, eds. 2016. *Springer Handbook of Robotics*. Springer Handbooks. Berlin: Springer.
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, et al. 2016. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529 (7587): 484–489.
- Siskind, Jeffrey Mark. 2001. 1975. "Grounding the Lexical Semantics of Verbs in Visual Perception Using Force Dynamics and Event Logic." *Journal of Artificial Intelligence Research* 15:31–90.
- Sünderhauf, Niko, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, et al. 2018. "The Limits and Potentials of Deep Learning for Robotics." *International Journal of Robotics Research* 37 (4–5): 405–420.
- Talmy, Leonard. 1988. "Force Dynamics in Language and Cognition." *Cognitive Science* 12 (1): 49–100.
- Tenorth, Moritz, and Michael Beetz. 2013. "KnowRob—a Knowledge Processing Infrastructure for Cognition-Enabled Robots." *International Journal of Robotics Research* 32 (5): 566–590. <http://ijr.sagepub.com/content/32/5/566.short>.
- Tenorth, Moritz, and Michael Beetz. 2015. "Representations for Robot Knowledge in the KnowRob Framework." *Artificial Intelligence*. San Diego: Elsevier.
- Thrun, Sebastian, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents. Cambridge, MA: MIT Press.
- Topp, Elin Anna, Maj Stenmark, Alexander Ganslandt, Andreas Svensson, Mathias Haage, and Jacek Malec. 2018. "Ontology-Based Knowledge Representation for Increased Skill Reusability in Industrial Robots." In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5672–5678. New York: IEEE. doi:10.1109/iroso.2018.8593566.
- Villani, Luigi, and Joris de Schutter. 2016. "Force Control." In Siciliano and Khatib 2016, 195–220. doi:10.1007/978-3-319-32552-1_9.
- Winston, Patrick Henry. 2012. "The Right Way." *Journal of Advances in Cognitive Systems* 1:23–36.
- Wolpert, Daniel. 2011. "The Real Reason for Brains." TED Talk. <https://www.youtube.com/watch?v=7s0cprfyyp8>.
- Wooldridge, Michael J. 2009. *An Introduction to Multiagent Systems*. 2nd ed. Hoboken, NJ: Wiley.
- Wörgötter, Florentin, Chris Geib, Minija Tamosiunaite, Eren Erdal Aksoy, Justus Piater, Hanchen Xiong, Ales Ude, et al. 2015. "Structural Bootstrapping—a Novel, Generative Mechanism for Faster and More Efficient Acquisition of Action-Knowledge." *IEEE Transactions on Autonomous Mental Development* 7 (2): 140–154.
- Wörgötter, Florentin, Eren Erdal Aksoy, Norbert Krüger, Justus Piater, Ales Ude, and Minija Tamosiunaite. 2013. "A Simple Ontology of Manipulation Actions Based on Hand-Object Relations." *IEEE Transactions on Autonomous Mental Development* 5 (2): 117–134.
- Yang, Yezhou, Ching L. Teo, Cornelia Fermüller, and Yiannis Aloimonos. 2013. "Robots with Language: Multi-label Visual Recognition Using NLP." In *2013 IEEE International Conference on Robotics and Automation*, 4256–4262. New York: IEEE.
- Zech, Philipp, Erwan Renaudo, Simon Haller, Xiang Zhang, and Justus Piater. 2019. "Action Representations in Robotics: A Taxonomy and Systematic Classification." *International Journal of Robotics Research* 38 (5): 518–562.
- Ziemke, Tom. 2001. "The Construction of 'Reality' in the Robot: Constructivist Perspectives on Situated Artificial Intelligence and Adaptive Robotics." *Journal of Foundations of Science* 6 (1): 163–233.

22 Abstract Concepts

Alessandro Di Nuovo

22.1 Introduction

One of the characteristics of human intelligence is the ability of thinking and reasoning about abstract concepts like “knowledge” and “beauty.” This ability is at the core of human innovation and creativity. In fact, it is required for fundamental capabilities such as the retrieval of past thoughts and memories, relational reasoning and problem-solving in current situations, and the processing of thoughts linked to the future (e.g., design, planning). Indeed, abstract concepts constitute an essential part of human language, where abstract words are often used in daily conversations to represent emotions, events, and situations that occur in physical environments and social interactions among people.

Human language includes concrete concepts, such as “water” or “glass,” that are linked to objects that can be objectively defined and understood. These are usually studied through a bottom-up approach that involves five major levels of analysis: phonetic, lexical, semantic, syntactic, and pragmatic. In contrast, abstract concepts like “love” or “freedom” don’t have specific physical referents; hence, they are more ambiguous, and their notion can significantly vary across individuals (Borghi et al. 2018). In this chapter, abstract concepts are broadly defined as higher-order, or complex, thoughts that are not bounded to a single, perceptually derived piece of information and that do not exist at any particular time or place (Barsalou 2003).

Even if the most common and intuitive definition of abstraction is opposite to that of concreteness, abstract and concrete concepts are not a dichotomy. They are considered part of a continuum (Barsalou and Wiemer-Hastings 2005), in which entities can have both abstract and concrete features in different proportions ranging from highly abstract (e.g., “justice”) to highly concrete (e.g., “stone”). The continuum view has gained strength in recent years, after growing evidence in support of embodied and grounded theories of cognition. In fact, a number of proposals have argued that abstract concepts can be grounded in a sensorimotor system as concrete concepts (see Pexman 2019) characterized by a continuum from unembodied (fully symbolic) to strongly embodied (Meteyard and Vigliocco 2008). A fundamental assumption of this view is that abstract concepts can be linked to embodied perceptions and learned through a process of progressive abstraction (Gentner and Asmuth 2019).

The embodied theories of the development of abstract thinking and reasoning constitute the theoretical resource for the design of artificial agents capable of abstract and symbolic processing, which is required for higher cognitive functions such as natural language understanding. This is one of the current challenges for the fast-growing field of cognitive robotics, in which future robots are expected to take on tasks once thought too complex or delicate to automate, especially in the fields of social care, companionship, therapy, domestic assistance, entertainment, and education (Matarić and Scassellati 2016; Di Nuovo et al. 2016).

This chapter aims at stimulating new research in cognitive robotics and artificial intelligence toward the creation of smarter robots that will be capable of understanding and manipulating abstract concept and words, thus overcoming the current limitations in human-robot communication by using natural language, which is the most intuitive of the user interfaces (Di Nuovo et al. 2018). To this end, section 22.2 provides a multidisciplinary background, briefly exploring recent embodied theories for the development of abstract concepts in humans. Section 22.3 will present pioneer work on cognitive robotics models of abstract words by implementing in robots the grounding transfer mechanism.

However, abstract concepts are not a single entity. They can be categorized into different domains that can be acquired using different strategies. Indeed, section 22.4 will present a different strategy for the embodied learning of numerical concepts that combines gestures and action with words, such as in the use of finger-counting representations to augment teaching a child (or a robot) about numbers. Numbers are a special domain of abstract concepts that constitute the building blocks of mathematics, a language of the human mind that can express the fundamental workings of the physical world and make the universe intelligible. Section 22.5 will present cognitive robotics models of emotion, another group that requires special attention among the abstract concepts since recent proposals that emotions can play an effective intermediary role for learning and grounding abstract concepts. Section 22.6 will discuss the current limitations in abstract cognition and robotics research. Finally, section 22.7 will give conclusions and identify future directions.

22.2 Education, Neuroscience, and Psychology Views on the Development of Abstract Concepts

Abstract concepts cover a vast domain, ranging from numbers to emotions and from social roles to mental state concepts. Anthropologists, cognitive scientists; developmental, social, and cognitive psychologists; educationalists; linguists; neuroscientists; and philosophers have extensively investigated how abstract concepts are acquired, used, and represented in the brain. This heterogeneity is one of the main reasons why it has been difficult to find a comprehensive theory that can account for the multiplicity of abstract concepts. This section will explore current views in education, neuroscience, and psychology characterized by an embodied approach to the development of abstract concepts.

The developmental psychologist Jean Piaget, whose work had an extensive influence on both theory and practice in education, argued that children develop abstract reasoning skills as part of their last stage of development, known as the formal operational stage, which usually occurs around the age of twelve (Piaget 1972). Specifically, this is the age at which most children transition from the concrete operational stage to the formal opera-

tional stage. However, brain-imaging studies have provided new evidence that there is a continuous neural development during adolescence that may last longer than what was theorized by Piaget. In particular, abstract reasoning requires maturational changes in some brain regions, such as the prefrontal cortex, which may last until late adolescence (Giedd and Rapoport 2010). Educational studies confirm that some tests of prefrontal lobe activity highly correlate with scientific reasoning ability and the capacity to reject scientific misconceptions and adopt correct ideas (Kwon and Lawson 2000). Other developmental psychologists (Harwood, Miller, and Vasta 2011) have argued that the development of abstract reasoning is not just a natural developmental stage; rather, it is the product of culture, experience, and teaching. Hayes and Kraemer (2017) explored cognitive neuroscience studies and presented evidence suggesting that sensorimotor processes can strengthen learning associated with the fundamental abstract concepts for understanding science, technology, engineering, and mathematics (STEM). On this basis, they proposed that embodied exercises could improve STEM pedagogy by situating abstract concepts in a concrete context, thus correlating intangible ideas with corporeal information. In doing so, rich multimodal distributed neural representations are forged, giving students a better chance at succeeding in the “hard” sciences, which are universally considered to be among the most abstract constructions of the human mind.

Numerous cognitive neuroscience studies suggested that both concrete and abstract concepts might be bodily grounded because they share similar mechanisms and modalities of representations, as both abstract and concrete concepts activate brain systems for action and perception (Gallese 2009). Behavioral and neurophysiological studies demonstrated a causal link between the motor system and the comprehension of both concrete and abstract language, where abstract concepts are acquired via a simulation process that calls on neural systems used in perceiving and acting on related concrete events (Glenberg et al. 2008). These results, also linked to the use of mirror neurons, support the embodied simulation theory (Gallese and Sinigaglia 2011), which provides a unitary explanation of basic abstract cognition, indicating that people reuse their own mental states or processes, represented in a bodily format, when functionally attributing them to others.

In the embodied cognition domain, at least three proposals have been offered to explain how abstract concepts could be acquired.

The first was proposed in the seminal work by Lakoff and Johnson (1980), who suggested that the meanings of abstract concepts could be grounded through conceptual metaphors (e.g., “love is a journey”), which help to embody abstract concepts into the sensorimotor experience. The linguistic and psychological evidence supporting the conceptual metaphors from the perspective of embodied simulations can be found in a review by Gibbs (2011). In this proposal, the evidence from the embodied cognition experiments should be explained in the light neural theory of thought and language; thus, he proposed that while children learn these metaphors, they develop conceptual metaphor neural circuits in connection to embodied experience, and these characterize abstract concepts. However, other authors (e.g., Murphy 1996; Dove 2011) criticized the developmental plausibility of this explanation, noting that children reach a mature metaphorical comprehension only quite late in middle childhood, at around ten years old. Several studies, however, show that metaphorical thinking emerges much earlier and constantly progresses, along with children’s knowledge and information-processing abilities (Vosniadou 1987). But it is not clear whether

these earlier developments in children's metaphorical thinking might contribute to the grounding of abstract concepts.

The second proposal assumes that the abstract concepts are mediated by language—that is, the conceptual grounding is augmented by concrete words (Dove 2014). In this context, the WAT (words as social tools) theory proposes a multiple representation view (Borghi et al. 2019), which attributes a major role to language and sociality in the acquisition of abstract concepts. Specifically, it hypothesizes that more abstract concepts are mainly linguistically acquired and induce in us a higher necessity to rely on others because of their complexity and our feelings of incompetence. Borghi et al. (2011) tested this idea in a study with adults, showing that learning novel abstract concepts was facilitated by verbal explanations (motor linguistic information) and not by manual actions, whereas the pattern was opposite for concrete concepts. By this view, the acquisition of language is a prerequisite for embodying abstract concepts. However, this proposal that abstract meaning is grounded through language is difficult to reconcile with strongly embodied developmental theories, like that of Glenberg and Gallese (2012), but it could be well associated with weak embodiment or hybrid models.

Howell, Jankowicz, and Becker (2005) suggested that children are likely to learn the first concrete words via direct experience. Later, abstract words are acquired, and their meanings are grounded by linguistic experience and by relationships to words learned earlier. According to Howell et al.'s model, children's representations of lexical cooccurrence information become increasingly sophisticated. Dove (2011) proposed a hybrid model in which language provides the child with new representational capacities (e.g., linguistic perceptual symbols) that support the learning of all kinds of concepts and are particularly helpful with characterizing abstract concepts.

Finally, a relatively recent idea is the proposal that abstract meaning is grounded through emotions (Vigliocco et al. 2013). The argument is that emotional experience should be considered a primary source of the embodied information that supports the development of abstract thinking and reasoning. Indeed, it forms a continuum that goes from sensorimotor experience that strongly characterizes concrete word representations to emotional experiences that dominate representations of abstract words (Moffat et al. 2015; Siakaluk, Knol, and Pexman 2014). Statistically, abstract words tend to have a stronger intensity of valence (good/bad, pleasant/unpleasant) than concrete words, making emotions an effective intermediary for learning and grounding abstract concepts (Altarriba, Bauer, and Benvenuto 1999). In this proposal, introspective emotion states could help the grounding of abstract meanings in embodied experience. Indeed, a significant step in forming abstract thinking occurs when, around two years of age, children start to learn words to express their emotions, mapping nonconcrete language to their felt experience for the first time. Kousta et al. (2011, 26) argued that “emotion may provide a bootstrapping mechanism for the acquisition of abstract words” because this process of learning labels for internal emotion states supports children in comprehending that words can identify entities that do not have an external, perceptual substantiation. Analyzing ratings of acquisition for abstract words by age, Kousta et al. (2011) showed that abstract words with a higher intensity of valence (e.g., “joy,” “grief”) were acquired earlier than neutral abstract words (e.g., “fashion,” “space”). Since emotional development continues throughout childhood, it seems likely that early grounding in emotion may be more about valence than about more complex emotions,

which develop later. However, the mechanism for the later acquisition of neutral abstract words is not fully explained by this proposal. Perhaps this might be facilitated through experiencing their use in the context of other words.

One of the current trends in the recent literature on abstract concepts focuses on the identification of the different domains and their corresponding brain representations (Borghi et al. 2017). In this respect, Desai, Reilly, and van Dam (2018) conducted a meta-analysis of the neural basis of four types of abstract concepts (numerical and emotional concepts and two higher-order abstract processes, morality judgments and theory of mind). Desai et al.'s (2018) analysis showed that the representation of abstract concepts is more widespread than is often assumed. Importantly, representations of different types of abstract concepts differ in important aspects, with each of the domains examined being associated with some unique areas of the brain. They found significant overlaps in the activation of morality and theory of mind concepts, which are likely processed when referring to social and episodic memories or to emotions and imagery. However, recent evidence suggests that defining concepts in terms of sole concreteness/abstractness is a simplification. Borghi et al. (2019) interviewed over three hundred adults and identified four domains of abstract concepts: philosophical-spiritual (e.g., sanctity), self-sociality (e.g., courtesy), emotive/inner states (e.g., anger), and physical, spatial, temporal, and quantitative (e.g., numbers).

Among the abstract domains, number concepts received special attention because of the strong relationship between the human mind and numerical cognition, which has made the latter a subject of research in the various disciplines that study the human mind and its development (Di Nuovo and Jay 2019). Their special role was confirmed by developmental, cross-cultural, and neuroscientific evidence that converges in the conclusion that number concepts occupy a range of positions on the continuum between abstract and concrete conceptual knowledge (Fischer and Shaki 2018). This includes the strong connection between spatial and mathematical domains (Young, Levine, and Mix 2018). Therefore, the study of numerical cognition can be a way to explore neuronal mechanisms of high-level brain functions (Nieder 2016). In fact, the observation of numerical practice within a situation can provide a provisional basis for pursuing the explanation of cognition as a nexus of relations between the mind at work and the world in which it works.

Number cognition is one of the skills that can be extended through embodied experiences from a rather limited set of inborn skills to an ever-growing network of abstract domains (Lakoff and Nuñez 2000). The early numerical practice is usually accompanied by gestures that are considered a window onto children's number knowledge because children spontaneously use gestures to convey information that is not necessarily found in their speech (Goldin-Meadow 1999). Within the human body, a special role is attributed to fingers, including a significant influence on the development of our system of counting. For example, we likely use a base-ten system because of the number of fingers we have. Indeed, recent research on the embodiment of mathematics has evidenced fingers as natural tools that play a fundamental role, from developing number sense to becoming proficient in basic arithmetic processing (Soylu, Lester, and Newman 2018).

These behavioral observations are confirmed by recent neuroimaging research in which empirical studies suggest there is a neural link or even a common substrate for the representation of numbers and fingers in the brain (for a review, see Peters and De Smedt 2018). Neuroimaging data show neural correlates of finger and number representations located

in neighboring or even overlapping cortex areas, suggesting that fingers may have a role in setting up the biological neural networks for more advanced (i.e., abstract) mathematical computations (Moeller et al. 2011). Importantly, several studies (e.g., Sato et al. 2007; Tschentscher et al. 2012) empirically showed the existence of a permanent neural link between the finger configurations and their cardinal number meaning in adults.

Emotions play a very important role in many aspects of our lives, including decision-making, perception, learning, and behavior, and emotional skills are an important component of human intelligence. The research on emotion concepts is intrinsically tied to the more general and controversial debate about the nature of emotion itself (Adolphs 2016). However, direct links between the body and the emotions have been long established. James (1894) provided the canonical example of such a link: “We know that we ‘fear’ a bear by perceiving changes in our own bodily state.” There is neuroscientific evidence that emotion changes the operating characteristics of cognition and action selection (Pessoa et al. 2019) and that there is, in fact, emotional activation before, during, or shortly after learning enhances memory (McGaugh 2018) and alters judgment (Gasper and Danube 2016). Given the importance of the body and its neural representation in emotion, it is perhaps unsurprising that the domain of emotion concepts has long been highlighted as a natural application for theories of embodied cognition. Indeed, almost all emotion theories consider that emotions are embodied via somatosensory, interoceptive, or motor information (Niedenthal and Ric 2017). Importantly, modern theories not only focus on embodiment but propose that emotions involve a cascade of events, with somatosensory and motor resources recruited at multiple time points in the perception, understanding, experience, and production of emotions (Winkielman, Coulson, and Niedenthal 2018).

22.3 Cognitive Robotics Models of Abstract Words

The design of cognitive robots that are capable of learning new words and concepts typically adopts an embodied and grounded approach. Chapter 20 introduced the “direct-grounding” approaches for developing language models in robots and presented applications of this strategy to learning more concrete words—that is, when the robot learns the names of objects it can perceive or words for actions it is performing or observing. For instance, robots can simulate the early stages of language development via the interaction of infants with caregivers (for a review, see Asada [2016]). Interestingly, Kawai et al. (2020) proposed a hidden Markov model to explain the development of syntactic categories that fit the developmental psychological experiments at different ages and for different languages.

The abstract/concrete continuum view of concepts suggests that the learning of higher-order, more abstract words may be obtained by extending the strategies and models for the grounding of concrete words. However, in the scientific literature only very few examples explore such an extension.

Recently, Cangelosi and Stramandinoli (2018) offered a review of two main strategies for grounding concepts without the sensorimotor experience of direct physical referents. In the “grounding-transfer” strategy (Cangelosi and Riga 2006), new concepts and words are learned by the robot in successive stages, via combining words whose meanings have been previously acquired through direct grounding. For example, a robot can learn the word “mermaid” if instructed to merge the previously acquired grounded meanings of “woman”

and “fish” and then transfer the result to the new word without ever seeing such a fantastic animal. In the alternative strategy, the robot learns abstract concepts by associating words to gestures and actions—for example, the use of finger counting to teach a child (or a robot) to count. In this section, we review some examples of the first strategy, while the second strategy is discussed in the next section, which presents cognitive robotics models of number cognition.

Recurrent neural networks (RNNs) are particularly suitable structures for modeling abstract concept learning since the recurrent connections allow the network to handle the sequence of progressive abstraction. Two main types of RNN were proposed: the Elman type, with a recursion on the hidden layer (Elman 1990), and the Jordan type, with a recursion from the output to the input (Jordan 1986).

From the “grounding transfer” view, Stramandinoli, Marocco, and Cangelosi (2012, 2017) investigated the problem of grounding intermediate abstract concepts—that is, higher-order actions that can be obtained by combining concrete motor concepts. Stramandinoli, Marocco, and Cangelosi (2012) performed experiments on a cognitive model for the humanoid robot iCub based on an RNN of the Elman-type, which permit the learning of higher-order concepts based on temporal sequences of action primitives and word sentences. The training of the model is incremental. The mechanism includes two stages: 1) the basic-grounding (BG) and 2) higher-grounding (HG) transfer mechanisms. During the BG, the robot learns a set of action primitives (e.g., “PUSH,” “GRASP” or “PULL,” “NEUTRAL”) using embodied and situated strategies. Two different stages were implemented for the HG training to enable different levels of the combination between basic and complex actions. In the first HG stage (i.e., HG-1), a sequence of previously learned words (e.g., “RECEIVE [is] PUSH [and] GRASP [and] PULL”) are provided to guide the hierarchical organization of the basic concepts directly grounded in sensorimotor experience (e.g., “PUSH,” “GRASP,” or “PULL”) in order to learn novel concepts (e.g., “GIVE”). Subsequently, the network receives as input the higher-order word “receive” and targets the outputs previously stored. During the second HG stage (i.e., HG-2), the robot learns three new higher-order words (“accept,” “reject,” “keep”) consisting of a combination of basic action primitives and higher-order words acquired during the previous HG-1 stage (e.g., “KEEP [is] PICK [and] NEUTRAL”). HG-2 adds a further hierarchical combination of words from both concrete concepts (BG) and the first level of abstraction words (HG-1). This training methodology is extremely flexible and permits designers to freely add novel words to the known vocabulary of the robot or to completely rearrange the word-meaning associations.

In follow-up work, Stramandinoli, Marocco, and Cangelosi (2017) proposed a partial RNN (Jordan-type) for learning the relationships between motor primitives and objects and performed experiments on the iCub robot for investigating the grounding of more abstract action words, such as “use” or “make.” Abstract action words represent a class of terms distant from the immediate perception that describe actions with a general meaning and that can refer to several events and situations. Therefore, they cannot be directly linked to sensorimotor experience through a one-to-one mapping with their physical referents in the world. The grounding of abstract action words is achieved through the integration of the linguistic, perceptual, and motor input modalities, recorded from the iCub sensors, in a three-layer RNN model (figure 22.1). The iCub robot first develops some basic perceptual and

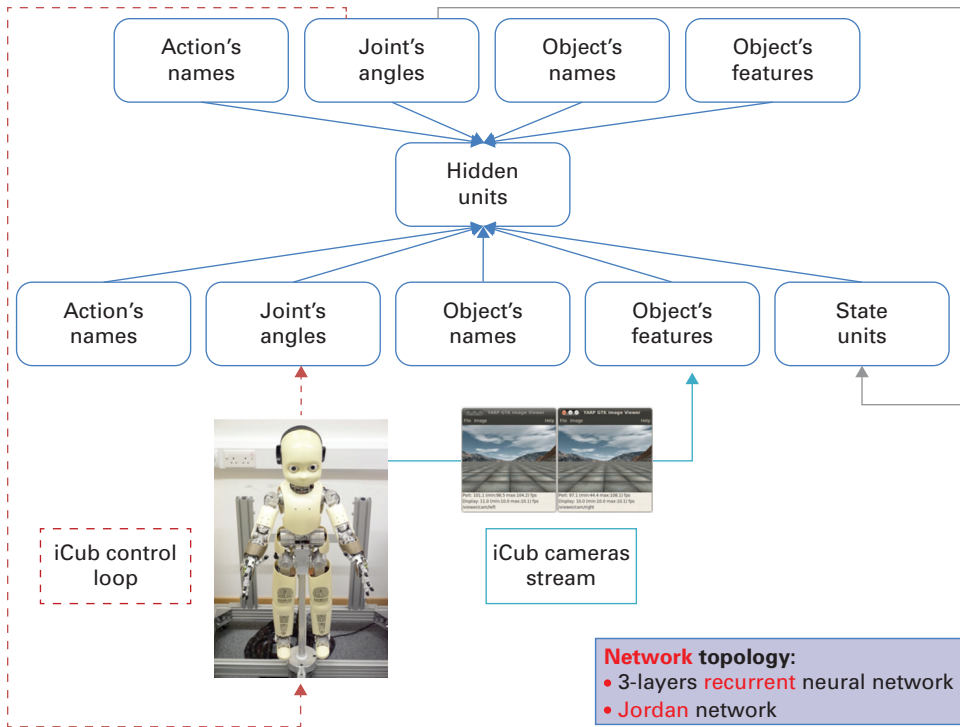


Figure 22.1
The partially recurrent neural network model for language abstraction.

motor skills, such as “PUSH,” “PULL,” and “LIFT,” necessary for initiating the physical interaction with the environment, and then it can use such knowledge to ground language. The training of the model is incremental and consists of three stages:

1. Prelinguistic—the robot is trained to recognize a set of objects (e.g., “KNIFE,” “HAMMER,” “BRUSH,” and so on) and learn object-related action primitives (e.g., “CUT,” “HIT,” “PAINT,” and so on) by combining low-level motor primitives. For example, the action primitive “cut” is built by iterating the “push-pull” sequence several times.

2. Linguistic-perceptual training—this is the first stage of language acquisition. The model is trained to associate labels with the corresponding object and actions (two-word sentences consisting of a verb followed by a noun—e.g., “CUT [with] KNIFE”). These words are directly grounded in perception and motor experience.

3. Linguistic abstract training—abstract action words (e.g., “USE,” “MAKE”) are grounded by combining and recalling the perceptual and motor knowledge previously linked to basic words (i.e., the previous linguistic-perceptual training). To derive the meaning of abstract action words, the robot, guided by linguistic instructions (e.g., “USE a KNIFE”), organizes the knowledge directly grounded in perception and motor knowledge. This phase of training represents the abstract stage of language acquisition when new concepts are formed by combining the meaning of terms acquired during the previous stages of training.

Novel lexical terms can be continually acquired throughout the robot’s development via new sensorimotor interactions with the environment that correspond to new linguistic

descriptions. At the end of the training, the robot was able to perform the behavior triggered by the linguistic description and the perceived object. The presence of clusters in the hidden units of the model suggested the formation of concepts from the multimodal data received as input by the network.

22.4 Cognitive Robotics Models of Numerical Concepts: Development and Representation

To explore embodied abstract cognition, cognitive robotics allows building embodied calculators that can merge abstract and concrete interpretations of numbers. This section concisely reviews some of the major computational models that were created to simulate the development of numerical cognition in artificial cognitive systems and robots. A more detailed review of the topic can be found in Di Nuovo and (Jay 2019).

In pure computational modeling, one of the milestones is the work of Ahmad, Casey, and Bale (2002), who introduced a very complex multinet network modular system following a mixture-of-experts approach. A peculiar aspect of the counting subsystem was a module for “pointing” to the next object to count “like a finger,” which was one of the first times that embodiment was included, even if its implications were not explicitly studied. The proposed architecture included two subsystems for subitizing and counting, which were realized by interconnecting several constituent modules, including connectionist networks that were trained independently. The main constituent architectures included, other than the multilayer feedforward neural network, recurrent connections of both Elman and Jordan types in the counting subsystem, and two self-organizing map (Kohonen 2001) architectures in the subitizing subsystem. The construction of this system also followed the assumption that subitizing is an innate capability, while counting should be learned via examples. This model has shown good adherence to the children’s data but also some inconsistency. For example, the simulation has a higher frequency of counting no objects compared to when children, who rarely make this error, count.

Chen and Verguts (2010) studied the interaction between the representations of number and space, presenting a bioinspired connectionist model that exhibited the SNARC effect in the parity judgment and number comparison tasks. The model was able to simulate not only the SNARC effect but also several other experimental data effects, including the spatial attention bias known as the Posner-SNARC effect and, after lesion, the spatial dysfunction found in patients with left-hemisphere damage. However, the “space representation” was hand-wired in such a way that it exhibited properties suggested by neuroscientific data.

The first attempt to use robots to explore embodied aspects of the interactions between numbers and space, made by Ruciński (2014), reproduced three psychological phenomena connected with number processing: size and distance effect, the SNARC effect, and the Posner-SNARC effect. The architecture was split into two neural pathways: “ventral,” which elaborates on the identity of objects and makes decisions according to the task and processes the language, and “dorsal,” which processes the spatial information—that is, locations and shapes of objects and sensorimotor transformations that provide direct support for visually guided motor actions. The results show that the embodied approach generated a more biologically plausible model by replacing arbitrary parts of the Chen

and Verguts model with elements that have direct physical connection and, therefore, more realistic interpretation.

In another experiment, Ruciński (2014) presented a new cognitive developmental robotics model to simulate aspects of the earlier work on gesture in counting by Alibali and DiRusso (1999), and indeed experimental results showed that pointing gestures significantly improved the counting accuracy of the humanoid robot iCub. The architecture was a recurrent neural network of the Elman type, with two input layers: one for the items to count—that is, a binary vector—and another for the proprioceptive information—that is, the arm and hand encoder values. The model was trained via backpropagation through time. Statistical analysis of the results showed adherence to the experimental data of Alibali and DiRusso.

Recently, Di Nuovo et al. conducted several experiments (De La Cruz et al. 2014; Di Nuovo, De La Cruz, and Cangelosi 2014; Di Nuovo et al. 2014) with the iCub humanoid robot to explore whether the association of finger counting with number words and/or visual digits could serve to bootstrap numerical cognition in a cognitive robot. The models (e.g., figure 22.2) were based on three RNNs of the Elman type, which were trained separately and then merged to learn the classification of the three inputs: finger counting (motor), digit recognition (visual), and number words (auditory)—that is, the triple-code




Figure 22.2

A schematic representation of the deep architecture for number cognition showing an integration of the models proposed by the several investigations of Di Nuovo et al. (2014).

model (Dehaene 1992). Also, the model mimics the two-hemisphere organization of the brain. Results of the various robotic experiments show that learning finger sequencing together with number word sequences speeds up the building of the neural network's internal links, resulting in a qualitatively better understanding (higher likelihood of the correct classification) of the real number representations.

Optimal cluster analysis (figure 22.3) showed that the internal representations of the finger configurations form the ideal basis for the building of an embodied number representation in the robot. Furthermore, it has been shown that such a cognitive developmental robotic model can subsequently sustain the robot's learning of the basic arithmetic operation of addition. However, this operation was implemented with an additional handcrafted layer just to show the possible further abstraction offered by the model.

Further investigation increased the biological adherence of the models and demonstrated the potential benefits, in terms of learning efficacy and efficiency, when used with deep-learning approaches, which are inspired by the complex layered organization and functioning of the cerebral cortex (Bengio 2009). Di Nuovo, De La Cruz, and Cangelosi (2015) created a model (e.g., figure 22.3) with an improved setup of the network weights employing restricted Boltzmann machines (RBMs) and the contrastive divergence-learning algorithm.

Follow-up studies (Di Nuovo 2017, 2018) focused on extending the simulation by incorporating the neural link observed between visual and motor areas in neuroscientific studies. Particularly, Di Nuovo (2018) investigated the long short-term memory architecture (Graves 2012) for learning to perform addition with the support of the robot's finger counting. Interestingly, the model showed similarities with studies with humans (children and adults) by performing an unusual number of split-five errors, which can be linked to the five finger representations (Domahs, Krinzinger, and Willmes 2008).

Di Nuovo and McClelland (2019) investigated the perceptual process of recognizing spoken digits in deep convolutional neural networks embodied in the iCub robot. Simulation results showed that the robot's fingers boost the performance by setting up the network and augmenting the training examples when these were numerically limited. This is a




Figure 22.3
Optimal leaf-order dendrogram of hidden-unit activation for finger configurations (*left*) and of number words (*right*).

common scenario in robotics, where robots will likely learn from a small amount of data. The embodied representation (finger encoder values) was compared to other representations, showing that fingers can represent the real counterpart of that artificial representation and can maximize learning performance. The results are associated with some behavior observed in several human studies in developmental psychology and neuroimaging. Overall, the hand-based representation provided our artificial system with information about magnitude representations that improved the creation of a more uniform number line, as seen in children (Gunderson, Spaepen, and Levine 2015). Importantly, this is the first time that a cognitive developmental robotics model has demonstrated effectiveness when compared against the standard approach for a benchmark machine-learning problem—that is, the Google Tensorflow Speech Recognition data set.

22.5 Cognitive Robotics Models of Emotions

The idea that robots may have emotions has captured the imagination of many researchers in the field of artificial intelligence, who have identified the crucial importance of emotions in the design of more intelligent and sociable robots (e.g., Breazeal 2004b; Fellous and Arbib 2005; Ziemke and Lowe 2009). The behavior-based robotic (BBR) has been a common approach for emotion-aware robots, which can use emotions as internal variables, which drive their external actions, mostly by correcting their operations according to the signals gained from their sensors (Arkin 2005). BBR ideas stimulated the design of robots capable of expressing emotional cues, such as the Kismet, Mexi, iCub, and Emys (Breazeal 2004a; Parmiggiani et al. 2012; Esau et al. 2003; Kędzierski et al. 2013). However, the mechanical expression of physical cues is just a preliminary step for the successful modeling of emotions; thus, emotionally capable cognitive architectures are necessary for enhancing the implementation of believable, autonomous, adaptive, and context-aware artificial agents (Hudlicka 2011).

Despite the theoretical agreement that the next generation of cognitive architectures must integrate emotion and cognition to define realistic models of human-machine interaction, in practice the computational modeling of emotion has been often underrated in cognitive architecture research. Models account for emotion as well as some other aspects of cognition, but usually, they are not aiming to be comprehensive architectures (see Rodríguez and Ramos 2015).

The computational modeling of emotion is frequently associated later with the addition of an emotion module that can influence some of the components of the general cognitive architecture (see Reizenzein et al. 2013). A notable example is SOAR (Laird 2012), which was not designed to model emotions; nevertheless, two different computational emotion models have been built upon SOAR: EMA (Marsella and Gratch 2009) and PEACTION (Marinier, Laird, and Lewis 2009). These two models represent the two principal alternative paths available to model emotions in cognitive architectures, and they also illustrate how theoretical assumptions in psychology can influence modeling choices. A general cognitive architecture designed to include emotions as flexible motivators for action is LIDA (Franklin et al. 2014), but this has only been considered at a conceptual level since modeling of emotions has not been implemented yet.

Pessoa (2017) identified two main categories of applications for emotion models in robotics: 1) to provide robots urgency to take action and make decisions, 2) to aid understanding of emotion in humans or to generate humanlike expressions. For the first category, significant applications of emotion-enabled general cognitive architectures have not yet been created for use with robots, even if general cognitive architectures have been used to control complex robots—for example, SOAR in the REEM robot (Puigbo et al. 2015). For the second category, it should be noted that many contributions in the robotics literature are loosely connected with the neuropsychological aspects of emotions, and the great majority fall under the category of pure machine-learning exercises, such as computer vision for facial expression recognition. Discussion and examples of recent contributions to modeling emotions in robotics can be found in the first volume of the book by Esposito and Jain (2016).

An example of the first category can be found in eMODUL, a perceptual system of emotion-cognition interaction specifically designed for robotics by Belkaid, Cuperlier, and Gaussier (2019). The eMODUL system is situated in its physical and social environment, and its components constantly appraise events from the body and the world, with a particular interest in emotionally relevant stimuli that affect other computational/cognitive processes (e.g., allocation of resources, organization of behavior). The system continuously processes emotionally modulated signals and reintegrates them into the information processing flow for higher-order processing. Valence extraction consists of the evaluation (appraisal) of the emotional values of complex representations. Therefore, the system sensations and actions are no longer neutral and objective but rather emotionally colored. For example, when occurring on the sensation space, emotional modulation affects perception and memory. When occurring on the action space, it can modulate action selection and motor expression. In terms of the system autonomy, these two types of modulations, respectively, have an impact on the allocation of cognitive/computational resources and the organization of appropriate behavior with regard to the system's survival, well-being, and task/goal demands. The authors provide two experimental examples of the application of the eMODUL system with artificial neural networks, in which emotional modulation consists of increasing or decreasing the synaptic efficacy of targeted populations of the neurons involved in these processes. The first experiment is in the context of a survival problem, in which a hunger modulation makes the robots more determined to access the resources and feed. The second is a visual search task designed similarly to the common experimental paradigm in psychology, in which the emotional (frustration or boredom) modulation of attention increases the robot's performance and fosters exploratory behavior to avoid deadlocks.

As an example of the second category, Prescott et al. (2019) included emotional signals in a neuroscience-inspired multimodal computational architecture for the autobiographical memory system, named the mental time travel model, to control the iCub robot. The model allows for retrieving past events, including their emotional associations, and projecting them into an imagined future by using the same system. This architecture proves useful for the social capabilities of robots by enabling face, voice (including emotion), action, and touch gesture recognition through interaction with humans. Using this system for imagining future events should allow for simulating and visualizing actions as well as planning actions before actual execution. This work is still at an early stage; however, experiments show that deploying emotionally mediated memory models into a brain-inspired

control architecture for the iCub robot has enhanced the robot's capability for recognizing social actors and actions.

22.6 Open Issues in Abstract Cognition and Robotics Research

In the interdisciplinary literature, most contributions recognize that to fully account for the representation of abstract concepts an extension beyond a purely grounded approach is needed. Pecher and Zeelenberg (2018) raised doubts on whether sensorimotor grounding alone can fully explain abstract concepts because recent evidence indicates that even concrete concepts are not always grounded in sensorimotor processes.

Another open issue has been highlighted by (Pexman 2019), who noted that so far none of the proposals for grounding abstract meaning have yet been tested in child studies. It will be important to investigate whether children's early abstract concepts are grounded through metaphor, language cooccurrence, and emotion. To this end, developmental robotics modeling can provide a powerful tool to collect preliminary information to evaluate or compare existing theories and to make novel experimental predictions that can be tested on humans (see chapter 3 for details). In particular, they could provide computational evidence in the debate on language development between "nativists" and "empiricists" (see chapter 20, section 1.1) by modeling the alternate theories and analyzing the resulting robot behavior in comparison to children's behavior.

To this end, computational models have the advantage of being fully specified in any implementation aspect, which makes them easily reproducible and verifiable, and they can produce detailed simulations of human performance in various situations and, for example, be used in experiments with any combination of stimuli. Furthermore, models can be lesioned (e.g., links between neurons can be cut) to simulate cognitive dysfunctions, and performance can be compared to the behavior of patients to gain information and insights into diagnosis and treatment that might be difficult to discover otherwise.

However, the cognitive robotics models proposed so far have been relatively naive because they focused on simulating only a particular aspect, verified with dummy tasks in simplified scenarios, and provided little evidence of their generalization ability in alternative, realistic settings. They considered only the concepts (e.g., metaphorical concepts such as "to grasp an idea") that have been empirically investigated in humans and found to be grounded in action and perception systems. Thus, we have yet to see if we might be able to extend these conclusions to other kinds of abstract concepts such as "politics" or "metaphysics." This is also the case with emotion modeling, which has predominantly been studied in terms of replicating human social behavior, while very little has been done to improve robots' abstract thinking. Significant improvement in the complexity of the models and, moreover, the test scenarios is needed before cognitive robotics modeling can be considered a reliable tool in education, neuroscience, and psychology research.

The reason for this lack of reality can be attributed not only to the limitations of current robotic platforms but also to the unavailability of raw data from children's experiments. Indeed, there are no open "benchmark" databases for cognitive robotics, unlike the typical open data behavior in machine learning. Robotic modelers can use only postprocessed data and statistical analyses for designing and validating models.

22.7 Conclusion

All these studies provided valuable information about the simulation of artificial learning and demonstrated the value of the cognitive robotics approach for studying aspects of abstract cognition. These findings reveal a novel way to achieve the humanization of artificial learning strategies, in which embodiment can make the robot's training more efficient and understandable for humans.

Further multidisciplinary research is required to gather data from children and get a better understanding of the underlying processes and strategies of abstract thinking and reasoning. It seems likely that there are developmental differences in the acquisition of the different types of concepts; therefore, hybrid models that combine sensorimotor experience and language appear to be viable options that should be investigated. In this respect, cognitive robotics can contribute to the theoretical development of abstract concepts acquisition and use in humans—that is, by providing a simulated environment for testing hypotheses—and benefit from the resulting discoveries to create innovative models of humanlike learning and social interaction.

To advance knowledge in this interdisciplinary field, we remark that closer collaboration among researchers in the multiple disciplines involved is necessary to share expertise and codesign studies. Importantly, we envision the need for real ad hoc joint experiments and for artificial simulations to obtain well-matched data comparing robots' and children's tasks. Furthermore, the availability of open databases will favor the engagement of the machine-learning community, as has occurred in other applied fields, such as computer vision, speech recognition, and DNA sequencing.

Additional Reading and Resources

- Book exploring the ways in which embodied and grounded cognition theories can be expanded into abstract words: Borghi, Anna, and Ferdinand Binkofski. 2014. *Words as Social Tools: An Embodied View on Abstract Concepts*. New York: Springer.
- This book presents a collection of studies that relate to various theoretical frameworks for abstract concepts, from neuroimaging to computational modeling and from behavioral experiments to corpus analyses: Bolognesi, Marianna, and Gerard Steen, eds. 2019. *Human Cognitive Processing, Vol. 65: Perspectives on Abstract Concepts: Cognition, Language and Communication*. Amsterdam: John Benjamins.
- Special issue with a collection of experimental and modeling papers on abstract concepts: Borghi, Anna M., Laura Barca, Ferdinand Binkofski, and Luca Tummolini. 2018. "Varieties of Abstract Concepts: Development, Use and Representation in the Brain." *Philosophical Transactions of the Royal Society B* 373 (1752): 20170121.
- Pearl, Lisa S., and Jon Sprouse. 2015. "Computational Modeling for Language Acquisition: A Tutorial with Syntactic Islands." *Journal of Speech, Language, and Hearing Research* 58 (3): 740–753.
- Source code and data for Di Nuovo and McClelland (2019): "Developing the Knowledge of Number Digits in a Child-Like Robot." *Nature Machine Intelligence* 1 (12): 594–605.

<http://doi.org/10.17032/shu-180017>. Number Understanding Modelling in Behavioral Embodied Robotic Systems (NUMBERS): <http://doi.org/10.17032/shu-180017>.

- Data set on concrete/abstract decision data for ten thousand English words in Pexman, P. M., et al. 2017. “The Calgary Semantic Decision Project: Concrete/Abstract Decision Data For 10,000 English Words.” *Behavior Research Methods* 49:407–417. <https://doi.org/10.3758/s13428-016-0720-6>.

References

- Adolphs, Ralph. 2016. “How Should Neuroscience Study Emotions? By Distinguishing Emotion States, Concepts, and Experiences.” *Social Cognitive and Affective Neuroscience* 12 (1): 24–31.
- Ahmad, Khurshid, Matthew Casey, and Tracey Bale. 2002. “Connectionist Simulation of Quantification Skills.” *Connection Science* 14 (3): 165–201.
- Alibali, Martha Wagner, and Alyssa DiRusso. 1999. “The Function of Gesture in Learning to Count: More than Keeping Track.” *Cognitive Development* 14 (1): 37–56.
- Altarriba, Jeanette, Lisa Bauer, and Claudia Benvenuto. 1999. “Concreteness, Context Availability, and Imageability Ratings and Word Associations for Abstract, Concrete, and Emotion Words.” *Behavior Research Methods, Instruments, and Computers* 31 (4): 578–602.
- Arkin, Ronald. 2005. “Moving Up the Food Chain: Motivation and Emotion in Behavior-Based Robots.” In *Who Needs Emotions? The Brain Meets the Robot*, 245–269. Series in Affective Science. Oxford: Oxford University Press.
- Asada, Minoru. 2016. “Modeling Early Vocal Development through Infant-Caregiver Interaction: A Review.” *IEEE Transactions on Cognitive and Developtmental Systems* 8 (2): 128–138.
- Barsalou, Lawrence. 2003. “Abstraction in Perceptual Symbol Systems.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 358:1177–1187.
- Barsalou, Lawrence, and Katja Wiemer-Hastings. 2005. “Situating Abstract Concepts.” In *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thought*, 129–163. Cambridge: Cambridge University Press.
- Belkaid, Marwen, Nicolas Cuperlier, and Philippe Gaussier. 2019. “Autonomous Cognitive Robots Need Emotional Modulations: Introducing the EMODUL Model.” *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 49 (1): 206–215.
- Bengio, Yoshua. 2009. *Learning Deep Architectures for AI: Foundations and Trends in Machine Learning*. Vol. 2. Netherlands: Now.
- Borghini, Anna, Laura Barca, Ferdinand Binkofski, Cristiano Castelfranchi, Giovanni Pezzulo, and Luca Tummolini. 2019. “Words as Social Tools: Language, Sociality and Inner Grounding in Abstract Concepts.” *Physics of Life Reviews* 29:120–153.
- Borghini, Anna, Laura Barca, Ferdinand Binkofski, and Luca Tummolini. 2018. “Varieties of Abstract Concepts: Development, Use and Representation in the Brain.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170121.
- Borghini, Anna, Ferdinand Binkofski, Cristiano Castelfranchi, Felice Cimatti, Claudia Scorolli, and Luca Tummolini. 2017. “The Challenge of Abstract Concepts.” *Psychological Bulletin* 143 (3): 263–292.
- Borghini, Anna, Andrea Flumini, Felice Cimatti, Davide Marocco, and Claudia Scorolli. 2011. “Manipulating Objects and Telling Words: A Study on Concrete and Abstract Words Acquisition.” *Frontiers in Psychology* 2:15.
- Breazeal, Cynthia. 2004a. “Function Meets Style: Insights from Emotion Theory Applied to HRI.” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 34 (2): 187–194.
- Breazeal, Cynthia. 2004b. “Social Interactions in HRI: The Robot View.” *Systems, Man, and Cybernetics, Part C* 34:181–186.
- Cangelosi, Angelo, and Thomas Riga. 2006. “An Embodied Model for Sensorimotor Grounding and Grounding Transfer: Experiments with Epigenetic Robots.” *Cognitive Science* 30 (4): 673–689.
- Cangelosi, Angelo, and Francesca Stramandinoli. 2018. “A Review of Abstract Concept Learning in Embodied Agents and Robots.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170131.
- Chen, Qi, and Tom Verguts. 2010. “Beyond the Mental Number Line: A Neural Network Model of Number-Space Interactions.” *Cognitive Psychology* 60 (3): 218–240.

- Dehaene, Stanislas. 1992. "Varieties of Numerical Abilities." *Cognition* 44:1–42.
- De La Cruz, Vivian, Alessandro Di Nuovo, Santo Di Nuovo, and Angelo Cangelosi. 2014. "Making Fingers and Words Count in a Cognitive Robot." *Frontiers in Behavioral Neuroscience* 8:13.
- Desai, Rutvik, Megan Reilly, and Wessel van Dam. 2018. "The Multifaceted Abstract Brain." *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170122.
- Di Nuovo, Alessandro. 2017. "An Embodied Model for Handwritten Digits Recognition in a Cognitive Robot." In *IEEE Symposium on Computational Intelligence, Cognitive Algorithms, Mind, and Brain (CCMB)*, 1–6. New York: IEEE.
- Di Nuovo, Alessandro. 2018. "Long-Short Term Memory Networks for Modelling Embodied Mathematical Cognition in Robots." In *Proceedings of the 2018 International Joint Conference on Neural Networks*, 1–7. New York: IEEE.
- Di Nuovo, Alessandro, Frank Broz, Filippo Cavallo, and Paolo Dario. 2016. "New Frontiers of Service Robotics for Active and Healthy Ageing." *International Journal of Social Robotics* 8 (3): 353–354.
- Di Nuovo, Alessandro, Frank Broz, Ning Wang, Tony Belpaeme, Angelo Cangelosi, Ray Jones, Raffaele Esposito, Filippo Cavallo, and Paolo Dario. 2018. "The Multi-modal Interface of Robot-Era Multi-robot Services Tailored for the Elderly." *Intelligent Service Robotics* 11 (1): 109–126.
- Di Nuovo, Alessandro, Vivian De La Cruz, and Angelo Cangelosi. 2014. "Grounding Fingers, Words and Numbers in a Cognitive Developmental Robot." In *IEEE Symposium on Cognitive Algorithms, Mind, and Brain (CCMB)*, 9–15. New York: IEEE.
- Di Nuovo, Alessandro, Vivian De La Cruz, and Angelo Cangelosi. 2015. "A Deep Learning Neural Network for Number Cognition: A Bi-cultural Study with the iCub." In *IEEE International Conference on Development and Learning and Epigenetic Robotics 2015*, 320–325. New York: IEEE.
- Di Nuovo, Alessandro, Vivian De La Cruz, Angelo Cangelosi, and Santo Di Nuovo. 2014. "The iCub Learns Numbers: An Embodied Cognition Study." In *International Joint Conference on Neural Networks*, 692–699. New York: IEEE.
- Di Nuovo, Alessandro, and Tim Jay. 2019. "Development of Numerical Cognition in Children and Artificial Systems: A Review of the Current Knowledge and Proposals for Multi-disciplinary Research." *Cognitive Computation and Systems* 1 (1): 2–11.
- Di Nuovo, Alessandro, and James L. McClelland. 2019. "Developing the Knowledge of Number Digits in a Child-Like Robot." *Nature Machine Intelligence* 1 (12): 594–605.
- Domahs, Frank, Helga Krinzinger, and Klaus Willmes. 2008. "Mind the Gap between Both Hands: Evidence for Internal Finger-Based Number Representations in Children's Mental Calculation." *Cortex* 44 (4): 359–367.
- Dove, Guy. 2011. "On the Need for Embodied and Dis-embodied Cognition." *Frontiers in Psychology* 1:242.
- Dove, Guy. 2014. "Thinking in Words: Language as an Embodied Medium of Thought." *Topics in Cognitive Science* 6 (3): 371–389.
- Elman, Jeffrey. 1990. "Finding Structure in Time." *Cognitive Science* 14 (2): 179–211.
- Esau, Natalia, Bernd Kleinjohann, Lisa Kleinjohann, and Dirk Stichling. 2003. "MEXI: Machine with Emotionally EXtended Intelligence." In *HIS*, 961–970. Amsterdam: IOS Press.
- Esposito, Anna, and Lakhmi Jain. 2016. *Toward Robotic Socially Believable Behaving Systems. Volume I: Modeling Emotions*. Berlin: Springer.
- Fellous, Jean-Marc, and Michael Arbib. 2005. *Who Needs Emotions? The Brain Meets the Robot*. Oxford: Oxford University Press.
- Fischer, Martin, and Samuel Shaki. 2018. "Number Concepts: Abstract and Embodied." *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170125.
- Franklin, Stan, Tamas Madl, Sidney D'Mello, and Javier Snider. 2014. "LIDA: A Systems-Level Architecture for Cognition, Emotion, and Learning." *IEEE Transactions on Autonomous Mental Development* 6 (1): 19–41.
- Gallese, Vittorio. 2009. "Motor Abstraction: A Neuroscientific Account of How Action Goals and Intentions Are Mapped and Understood." *Psychological Research PRPF* 73 (4): 486–498.
- Gallese, Vittorio, and Corrado Sinigaglia. 2011. "What Is so Special about Embodied Simulation?" *Trends in Cognitive Sciences* 15 (11): 512–519.
- Gasper, Karen, and Cinnamon Danube. 2016. "The Scope of Our Affective Influences: When and How Naturally Occurring Positive, Negative, and Neutral Affects Alter Judgment." *Personality and Social Psychology Bulletin* 42 (3): 385–399.
- Gentner, Dedre, and Jennifer Asmuth. 2019. "Metaphoric Extension, Relational Categories, and Abstraction." *Language, Cognition and Neuroscience* 34 (10): 1298–1307.
- Gibbs, Raymond. 2011. "Evaluating Conceptual Metaphor Theory." *Discourse Processes* 48 (8): 529–562.

- Giedd, Jay, and Judith Rapoport. 2010. "Structural MRI of Pediatric Brain Development: What Have We Learned and Where Are We Going?" *Neuron* 67 (5): 728–734.
- Glenberg, Arthur, and Vittorio Gallese. 2012. "Action-Based Language: A Theory of Language Acquisition, Comprehension, and Production." *Cortex* 48 (7): 905–922.
- Glenberg, Arthur, Marc Sato, Luigi Cattaneo, Lucia Riggio, Daniele Palumbo, and Giovanni Buccino. 2008. "Processing Abstract Language Modulates Motor System Activity." *Quarterly Journal of Experimental Psychology* 61 (6): 905–919.
- Goldin-Meadow, Susan. 1999. "The Role of Gesture in Communication and Thinking." *Trends in Cognitive Sciences* 3 (11): 419–429.
- Graves, Alex. 2012. "Long Short-Term Memory." In *Supervised Sequence Labelling with Recurrent Neural Networks: Studies in Computational Intelligence*, edited by Alex Graves, 37–45. Berlin: Springer.
- Gunderson, Elizabeth, Elizabeth Spaepen, and Susan Levine. 2015. "Approximate Number Word Knowledge before the Cardinal Principle." *Journal of Experimental Child Psychology* 130:35–55.
- Harwood, Robin, Scott Miller, and Ross Vasta. 2011. *Child Psychology: Development in a Changing Society*. 5th ed. Hoboken, NJ: John Wiley and Sons.
- Hayes, Justin, and David Kraemer. 2017. "Grounded Understanding of Abstract Concepts: The Case of STEM Learning." *Cognitive Research: Principles and Implications* 2 (1): 7.
- Howell, Steve, Damian Jankowicz, and Suzanna Becker. 2005. "A Model of Grounded Language Acquisition: Sensorimotor Features Improve Lexical and Grammatical Learning." *Journal of Memory and Language* 53 (2): 258–276.
- Hudlicka, Eva. 2011. "Guidelines for Designing Computational Models of Emotions." *International Journal of Synthetic Emotions* 2 (1): 26–79.
- James, William. 1894. "The Physical Basis of Emotion." *Psychological Review* 1 (2): 516–529.
- Jordan, Michael. 1986. "Attractor Dynamics and Parallelism in a Connectionist Sequential Machine." In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, 531–546. Amherst, MA: Erlbaum Associates.
- Kawai, Yuji, Yuji Oshima, Yuki Sasamoto, Yukie Nagai, and Minoru Asada. 2020. "A Computational Model for Child Inferences of Word Meanings via Syntactic Categories for Different Ages and Languages." *IEEE Transactions on Cognitive and Developmental Systems* 12 (3): 401–416.
- Kędzierski, Jan, Robert Muszyński, Carsten Zoll, Adam Oleksy, and Mirela Frontkiewicz. 2013. "EMYS—Emotive Head of a Social Robot." *International Journal of Social Robotics* 5 (2): 237–249.
- Kohonen, Teuvo. 2001. *Self-Organizing Maps*. Berlin: Springer.
- Kousta, Stavroula-Thaleia, Gabriella Vigliocco, David Vinson, Mark Andrews, and Elena Del Campo. 2011. "The Representation of Abstract Words: Why Emotion Matters." *Journal of Experimental Psychology: General* 140 (1): 14–34.
- Kwon, Yong-Ju, and Anton Lawson. 2000. "Linking Brain Growth with the Development of Scientific Reasoning Ability and Conceptual Change during Adolescence." *Journal of Research in Science Teaching* 37 (1): 44–62.
- Laird, John. 2012. *The SOAR Cognitive Architecture*. Cambridge, MA: MIT Press.
- Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Lakoff, George, and Rafael Nuñez. 2000. *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.
- Marinier, Robert, John Laird, and Richard Lewis. 2009. "A Computational Unification of Cognitive Behavior and Emotion." *Cognitive Systems Research* 10 (1): 48–69.
- Marsella, Stacy, and Jonathan Gratch. 2009. "EMA: A Process Model of Appraisal Dynamics." *Cognitive Systems Research* 10 (1): 70–90.
- Matarić, Maja, and Brian Scassellati. 2016. "Socially Assistive Robotics." In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1973–1994. Cham, Switzerland: Springer.
- McGaugh, James. 2018. "Emotional Arousal Regulation of Memory Consolidation." *Current Opinion in Behavioral Sciences* 19:55–60.
- Meteyard, Lotte, and Gabriella Vigliocco. 2008. "The Role of Sensory and Motor Information in Semantic Representation: A Review." In *Perspectives on Cognitive Science*, edited by Paco Calvo and Antoni Gomila, 291–312. San Diego: Elsevier.
- Moeller, Korbinian, Laura Martignon, Silvia Wessolowski, Joachim Engel, and Hans-Christoph Nuerk. 2011. "Effects of Finger Counting on Numerical Development—the Opposing Views of Neurocognition and Mathematics Education." *Frontiers in Psychology* 2:328.

- Moffat, Michael, Paul Siakaluk, David Sidhu, and Penny Pexman. 2015. "Situated Conceptualization and Semantic Processing: Effects of Emotional Experience and Context Availability in Semantic Categorization and Naming Tasks." *Psychonomic Bulletin and Review* 22 (2): 408–419.
- Murphy, Gregory. 1996. "On Metaphoric Representation." *Cognition* 60 (2): 173–204.
- Niedenthal, Paula, and François Ric. 2017. *Psychology of Emotion*. Hove, UK: Psychology Press.
- Nieder, Andreas. 2016. "The Neuronal Code for Number." *Nature Reviews Neuroscience* 17:366.
- Parmiggiani, Alberto, Marco Maggiali, Lorenzo Natale, Francesco Nori, Alexander Schmitz, Nikos Tsagarakis, José Santos Victor, Francesco Becchi, Giulio Sandini, and Giorgio Metta. 2012. "The Design of the iCub Humanoid Robot." *International Journal of Humanoid Robotics* 09 (4): 1250027.
- Pecher, Diane, and René Zeelenberg. 2018. "Boundaries to Grounding Abstract Concepts." *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170132.
- Pessoa, Luiz. 2017. "Do Intelligent Robots Need Emotion?" *Trends in Cognitive Sciences* 21 (11): 817–819.
- Pessoa, Luiz, Loretta Medina, Patrick Hof, and Ester Desfilis. 2019. "Neural Architecture of the Vertebrate Brain: Implications for the Interaction between Emotion and Cognition." *Neuroscience and Biobehavioral Reviews* 107:296–312.
- Peters, Lien, and Bert De Smedt. 2018. "Arithmetic in the Developing Brain: A Review of Brain Imaging Studies." *Developmental Cognitive Neuroscience* 30:265–279.
- Pexman, Penny. 2019. "The Role of Embodiment in Conceptual Development." *Language, Cognition and Neuroscience* 34 (10): 1274–1283.
- Piaget, J. 1972. "Intellectual Evolution from Adolescence to Adulthood." *Human Development* 15 (1): 1–12.
- Prescott, Tony, Daniel Camilleri, Uriel Martinez-Hernandez, Andreas Damianou, and Neil Lawrence. 2019. "Memory and Mental Time Travel in Humans and Social Robots." *Philosophical Transactions of the Royal Society B: Biological Sciences* 374 (1771): 20180025.
- Puigbo, Jordi-Ysard, Albert Pumarola, Cecilio Angulo, and Ricardo Tellez. 2015. "Using a Cognitive Architecture for General Purpose Service Robot Control." *Connection Science* 27 (2): 105–117.
- Reisenzein, Rainer, Eva Hudlicka, Mehdi Dastani, Jonathan Gratch, Koen Hindriks, Emiliano Lorini, and John-Jules Meyer. 2013. "Computational Modeling of Emotion: Toward Improving the Inter- and Intradisciplinary Exchange." *IEEE Transactions on Affective Computing* 4 (3): 246–266.
- Rodríguez, Luis-Felipe, and Félix Ramos. 2015. "Computational Models of Emotions for Autonomous Agents: Major Challenges." *Artificial Intelligence Review* 43 (3): 437–465.
- Ruciński, Marek. 2014. "Modelling Learning to Count in Humanoid Robots." PhD thesis, University of Plymouth, UK.
- Sato, Marc, Luigi Cattaneo, Giacomo Rizzolatti, and Vittorio Gallese. 2007. "Numbers within Our Hands: Modulation of Corticospinal Excitability of Hand Muscles during Numerical Judgment." *Journal of Cognitive Neuroscience* 19:684–693.
- Siakaluk, Paul, Nathan Knol, and Penny Pexman. 2014. "Effects of Emotional Experience for Abstract Words in the Stroop Task." *Cognitive Science* 388 (8): 1698–1717.
- Soylu, Firat, Frank Lester Jr., and Sharlene D. Newman. 2018. "You Can Count on Your Fingers: The Role of Fingers in Early Mathematical Development." *Journal of Numerical Cognition* 4 (1): 107–135.
- Stramandinoli, Francesca, Davide Marocco, and Angelo Cangelosi. 2012. "The Grounding of Higher Order Concepts in Action and Language: A Cognitive Robotics Model." *Neural Networks* 32:165–173.
- Stramandinoli, Francesca, Davide Marocco, and Angelo Cangelosi. 2017. "Making Sense of Words: A Robotic Model for Language Abstraction." *Autonomous Robots* 41 (2): 367–383.
- Tschentscher, Nadja, Olaf Hauk, Martin Fischer, and Friedemann Pulvermüller. 2012. "You Can Count on the Motor Cortex: Finger Counting Habits Modulate Motor Cortex Activation Evoked by Numbers." *NeuroImage* 594 (4): 3139–3148.
- Vigliocco, Gabriella, Stavroula-Thaleia Kousta, Pasquale Anthony Della Rosa, David Vinson, Marco Tettamanti, Joseph Devlin, and Stefano Cappa. 2013. "The Neural Representation of Abstract Words: The Role of Emotion." *Cerebral Cortex* 24 (7): 1767–1777.
- Vosniadou, Stella. 1987. "Children and Metaphors." *Child Development* 58 (3): 870–885.
- Winkelman, Piotr, Seana Coulson, and Paula Niedenthal. 2018. "Dynamic Grounding of Emotion Concepts." *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170127.
- Young, Christopher, Susan Levine, and Kelly Mix. 2018. "The Connection between Spatial and Mathematical Ability across Development." *Frontiers in Psychology* 9:755.
- Ziemke, Tom, and Robert Lowe. 2009. "On the Role of Emotion in Embodied Cognitive Architectures: From Organisms to Robots." *Cognitive Computation* 1 (1): 104–117.

23 Robots and Machine Consciousness

Antonio Chella

23.1 Introduction

Building a conscious robot is an enormous scientific and technological challenge. Debates about the possibility of sentient robots and the positive outcomes and risks for human beings are no longer confined to philosophical circles. Consciousness is part of the physical world, and therefore its aspects can be studied and even replicated by robot systems.

There is no accepted definition of consciousness so far. Searle (2000) claimed that “consciousness consists of inner, qualitative, subjective states and processes of sentience or awareness. Consciousness, so defined, begins when we wake in the morning from a dreamless sleep and continues until we fall asleep again, die, go into a coma, or otherwise become ‘unconscious’” (559). Vimal (2009) overviewed several meanings of the word employed in scientific works related to the study of consciousness.

Although there are contrasting philosophical positions concerning consciousness (see, e.g., Blackmore and Troscianko [2018] for an up-to-date review), it is useful to point out the broad distinction of consciousness as *experience* versus consciousness as *function*. For experience, a subject is conscious when they feel visual experiences, bodily sensations, mental images, and emotions (Chalmers 1995). As Nagel (1974) pointed out, a subject has a conscious experience if there is something that is like to be that subject.

For function, a conscious subject can integrate information (Tononi 2008); they process information that is globally available (Dehaene et al. 2017); they are introspectively aware of themselves (Floridi 2005). Moreover, they possess an inner model of themselves and of the external environment (Holland 2003b). They can anticipate perceptual and behavioral activities (Hesslow 2002). They generate inner speech (Morin 2005) and act by sensorimotor interactions with the external world (O’Regan and Noë 2001), among other capabilities.

In brief, the multidisciplinary effort of robot and machine consciousness is aimed at investigating consciousness in the light of robotics and artificial systems, psychology, philosophy of mind, ethics, and neuroscience. The broad scopes of robot and machine consciousness are:

- to build robots that show forms of functional consciousness by taking inspiration from biological consciousness;
- to build robots based on theoretical issues of consciousness;

- to employ robots as tools to model and to understand biological aspects of consciousness;
- to study procedures aimed at measuring consciousness in robots;
- to discuss ethical problems emerging through the overlap of robotics and consciousness.

23.2 A Brief History of Robot Consciousness

To the best of the author's knowledge, the first occurrence of the word "artificial consciousness" is found in the book *Cybernetic Machines* by T. N. Nemes, published in Hungary in 1962. The book was translated into English in 1970. Nemes, in this early attempt, considered artificial consciousness as the capability of a robot to discriminate between self and others. The author proposed a conceptual sketch of a circuit able to distinguish between proprioceptive inputs that generate sentences as "I go" from shape recognition and motion perception circuits that process data from external inputs able to create sentences as "Peter goes."

The modern scientific framework of artificial and robot consciousness has been primarily introduced by Igor Aleksander (1992, 2015). At the ICANN 1992 Conference in Brighton, Aleksander presented a paper on capturing consciousness in neural systems, where he proposed the postulates defining a conscious organism that may be applied to a biological organism or an artifact. Notably, during the invited talk, Aleksander announced that the "hunting season of artificial consciousness is open."

Another influential early model for machine consciousness is due to Schmidhuber (1992). He discussed machine consciousness by presenting an unsupervised neural network able to discover and learn unexpected events.

The symposium on "Can a Machine Be Conscious," organized by the Swartz Foundation in 2001, was another milestone for robot consciousness. The concluding remarks of Christof Koch, valid still today, stated that "we know of no fundamental law or principle operating in this universe that forbids the existence of subjective feelings in artifacts designed or evolved by humans."¹

Since 2001, many conferences, workshops, and special issues of journals have been devoted to the field of robot consciousness. Early works are described in the collections edited by Holland (2003a), Clowes et al. (2007), and Chella and Manzotti (2007b). In 2007, the Association for the Advancement of Artificial Intelligence (AAAI) organized a fall symposium on "AI and Consciousness," with the proceedings edited by Chella and Manzotti (2007a).

Reggia (2013) provided quite an up-to-date review of the field. A collection of recent research papers concerning consciousness in humanoid robots was edited by Chella et al. (2019).

During the summer of 2017, SRI International organized a series of workshops on technology and consciousness. The workshops provided a general view of machine consciousness; the outcomes are summarized in a technical report edited by Rushby and Sanchez (2018).

A continuous source of information is the *Journal of Artificial Intelligence and Consciousness (JAIC)*, formerly known as the *International Journal of Machine Consciousness* and edited by World Scientific Press.

23.3 Robot Consciousness and Neuroscience

Consciousness is an important research topic in neuroscience (Rees et al. 2002; Tononi and Koch 2008; Koch et al. 2016). Many neuroscientists working on consciousness have built computational models to test their theories.

The late Nobel Prize winner Gerald Edelman, a scholar of research on biological consciousness, employed robots to validate parts of his theory. Reeke et al. (1990) discussed the Darwin series of automata (see chapter 1 for their influence in the history of cognitive robotics). They are computational systems that incorporate models of synaptic modifications, of the organization of neural cells in large assemblies, and of the integration of the actions of different cortical layers to generate the behavior of a robot according to context and its history and without the need for preprogramming the robot. Darwin I is a simple network able to recognize patterns, while Darwin II can categorize and generate associations. Darwin III is a sophisticated robot model working in a simulated environment and able to learn sensorimotor coordination, the capability of tracking objects, and the ability to reach and grasp objects and to categorize them by interacting with the environment.

Krichmar et al. (2005) discussed complex systems implemented on a real moving robot and based on computational simulations of parts of the nervous system. Darwin VII can carry out perceptual categorization and conditioned responses in simple foraging tasks, and Darwin VIII can solve the binding problem—that is, to bind the attributes of a perceived scene to form suitable coherent categories, without the need of a control system. The robot behavior emerges from the interaction of different cell assemblies without the need for preprogramming.

Stanislas Dehaene, a world-leading expert on biological consciousness, built several computational models of the neural correlates of consciousness (Dehaene et al. 2003; Zylberberg et al. 2010). In more detail, Dehaene et al. (2003) describe a computational model based on two spaces. The first space is a global neural workspace made up of distributed neurons tightly interconnected with long-range axons. The second space is a set of specialized processors related to perception, motion, memory, attention, and evaluation. Briefly, the role of the first space is to broadcast the information coming from the specialized processors belonging to the second space. The global neural workspace is tightly related to the global workspace theory (see below).

Paul Verschure (2013) analyzed the core principles of conscious states and proposed a biologically inspired architecture for perception, cognition, and action (DAC, or distributed adaptive control) to implement the core principles. Verschure claimed that the shift of research from artificial intelligence to artificial consciousness would bring more advanced machines and address the critical problem of subjective experience in humans and machines.

Recently, Dehaene et al. (2017) discussed the possibility of machine consciousness in the prestigious journal *Science*. They proposed a separation of two different information-processing aspects related to consciousness. The first aspect is related to the selection of information for global broadcasting. A second aspect is correlated to self-monitoring of these computations. The article reviewed examples of computational models inspired to machine consciousness, and it concluded with the claim that “the empirical evidence is compatible with the possibility that consciousness arises from nothing more than specific computations” (Dehaene et al. 2017, 7).

23.4 Theoretical Issues of Consciousness in Humans and Robots

A common route of investigation in robot and machine consciousness is to find a minimal set of characteristics that should be verified in an artifact before asserting whether the artifact is conscious or not.

Aleksander (1992), in the previously cited attempt, proposed five axioms that should be verified by a conscious organism. They are as follows: 1) an organism that does not learn cannot be conscious; 2) a conscious organism possesses an inner state able to represent the external world; 3) a conscious organism is able to pay attention to the contents of its internal state; 4) a conscious organism is able to generate inner states related to sequences of external inputs and to generate suitable actions; 5) the organism is able to predict external events by controlled developments of its inner state.

Aleksander and Dunmall (2003) extended this early attempt and proposed a new set of axioms for minimal consciousness in agents. These axioms are the minimal mechanisms underpinning experience. It should be noted that these authors are interested in finding a theoretical grounding for experiential consciousness in humans and artifacts. The axioms are derived from the introspective analysis of consciousness.

Let A be a generic agent in the world S . For A to be conscious of S :

- A has perceptual states that represent parts of S , corresponding to the subjective feeling that the conscious subject A is a part of, but separate from, the world S ;
- A has internal states that recall elements of S or generate imagined S -like sensations, corresponding to the subjective feeling that the perception of the world S is mixed with A 's past experiences;
- A can pay attention to parts of S to represent or to imagine, corresponding to the reflective feeling that A 's experience of the world S is selective;
- A can control imagined state sequences to generate a plan of action, corresponding to the reflective feeling that A can think ahead of time to decide what to do;
- A has affective states able to evaluate planned operations and determine the appropriate action, corresponding to the subjective feeling that A has emotions and moods that determine its course of activities.

Aleksander and Dunmall translated these axioms in terms of mathematical constraints to be satisfied by a neural system to be considered as endowed with minimal consciousness. Aleksander (2005) proposed a schema of a cognitive architecture derived from the axioms (figure 23.1).

Selmer Bringsjord (see, e.g., Bringsjord 2007) contrasted the possibility of experiences in robots and proposed the notion of *cognitive* consciousness defined in terms of formal axioms of deontic cognitive event calculus (DCEC*; Bringsjord et al. 2018). DCEC* is a logical framework based on multisorted, quantified modal logic. It considers operators for belief, intention, knowledge, obligation, and so on. The framework allows the representation of formulae for belief and obligation. It is a family of logic in which the personal pronoun I^* is based on provable theorems.

The framework provided by Bringsjord and colleagues considers the cognitive aspects of consciousness because it represents the belief about oneself and is related to a first-

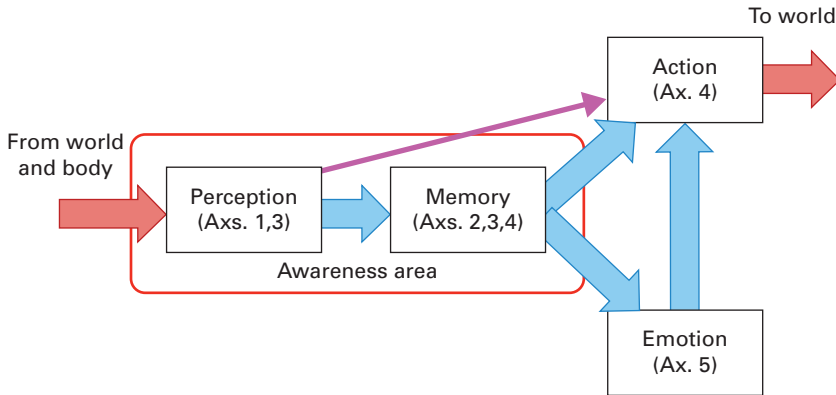


Figure 23.1

The cognitive architecture proposed by Aleksander (2005) summarizing the consciousness axioms by Aleksander and Dunmall (2003).

person representation of self-consciousness, but without considering bodily experiences. HyperSlate™ is a freely available implementation of the framework (see link in the list of additional resources).

Bringsjord et al. (2015) reported an impressive example of the framework by presenting an implementation on the NAO robot that passed the human test of self-consciousness proposed by Floridi (2005).

Giulio Tononi proposed the information integration theory (IIT) of consciousness. IIT is today the most debated scientific theory of consciousness, and many scholars actively contribute to the theory. Important outcomes also follow for robot consciousness.

The original formulation (Tononi and Sporns 2003; Tononi 2004; Tononi 2008) starts from the observation that conscious experience is differentiated because the potential repertoire of different conscious states is enormous. At the same time, conscious experience is integrated, as every conscious state is experienced as a single entity. Thus, the substrate of conscious experience must be an integrated entity able to differentiate among an enormous repertoire of different states.

The capability of a system S to differentiate among states is related to how much information can be generated by the system, and it is measured by the entropy of the system $H = -\sum p_i \log_2 p_i$, where p_i are the probabilities of the alternative outcomes of the system S .

The capability of a system S to integrate information can be measured through the effective information EI . Let us consider the system S subdivided into two partitions $[A, B]$, and let us perturb A in order to reach the maximum entropy to outputs of A —that is, $A^{H_{max}}$. Then, the effective information from A to B is given by $EI(A \rightarrow B) = MI(A^{H_{max}}, B)$, where $MI(A, B) = H(A) + H(B) - H(AB)$ is the mutual information that measures the information shared by the source A and the target B .

The effective information EI is a measure of how the subsystem B is connected with the subsystem A . Let us consider the system $S1$ in figure 23.2 (top), where there are tight connections from A to B . Then, when A is highly perturbed, B will produce many different outputs, and $EI(A \rightarrow B)$ will be a high value.

Instead, if there are scarce or low connections between A and B, as in the case of system S2 in figure 23.2 (*bottom*), then the perturbation of A will produce scarce effects on B, and thus $EI(A \rightarrow B)$ will be a small or null value. The effective information is generally nonsymmetric, so, for a given partition, the effective information is the sum of the EI for both directions: $EI(A \leftrightarrow B) = EI(A \rightarrow B) + EI(B \rightarrow A)$. It is to be noted that if there is a partition $[A, B]$ of the system S so that $EI(A \leftrightarrow B) = 0$, then S is made up by the two independent subsystems A and B.

To measure the capability of the system to integrate information, we need to find the minimum information bipartition $MIB(S) = [A, B]$ —that is, the partition $[A, B]$ of the system S for which the normalized effective information leads a minimum. $\Phi(S)$ measures the capability of the system S to integrate information, and it is the effective information given by the minimum information partition: $\Phi(S) = EI(MIB(S))$.

A subset of the system S with $\Phi > 0$ is called a *complex* when it is not included within a more substantial subset of S with a higher value of Φ . The complex of the system S with the maximum amount of $\Phi(S)$ is the *main complex*. Tononi (2004) claims that the main complex contributes to the conscious experience of S, and the measure $\Phi(S)$ grades the consciousness of the system.

Therefore, a conscious complex is a complex with a high value of $\Phi(S)$. The other parts of the systems do not contribute to the consciousness of the system. He supports his claim by analyzing different neural network models of parts of the brain and by showing that the networks with high values of $\Phi(S)$ are those typically associated with consciousness.




Figure 23.2

A pictorial view of a system subdivided into connected partitions A and B. *Top*: The two partitions of S1 are tightly connected, and $EI(A \rightarrow B)$ will be a high value. *Bottom*: The two partitions of S2 are barely connected, and $EI(A \rightarrow B)$ will be a low value.

Koch (2009) indicates some of the challenges of the IIT to be the unclear relationship of high values of $\Phi(S)$ with intelligence, the need for efficient algorithms for computing $\Phi(S)$ in real systems, and the need to clarify the relationships between conscious and unconscious processing.

It is to be noted that the original $\Phi(S)$ is a static measure of S ; that is, it depends on the connections of the subparts of S and not on its dynamics. Balduzzi and Tononi (2008) generalize the IIT by considering the dynamics of the system. Several other extensions of IIT have been proposed in the literature; the most up-to-date version is in Oizumi et al. (2014). Tegmark (2016) investigates many variants of the original $\Phi(S)$ measure to derive exact and approximated versions that are computationally feasible to apply to real-world data.

According to IIT, experience—for example, information integration—is a fundamental quantity of nature as the mass, the charge, and the energy. Any physical system may have experiences to the extent that it can integrate information. Therefore, it could be possible in principle to build conscious artifacts by endowing them with a complex of high $\Phi(S)$. However, Koch and Tononi (2017) suggest that conventional computer architectures are unable to perform an effective integration of information, and they are unable to experience anything. A robot based on a conventional computer may be a “zombie,” an entity similar to a conscious entity from its outside behavior but incapable of having real experience. Unconventional architectures, such as the neuromorphic systems, are more likely to perform the effective information integration processes happening in the brain, and therefore, they are more likely to have experience.

According to the analysis of Koch and Tononi (2008), there are many unessential ingredients for consciousness, in the sense that they have no roles in information integration. Sensory inputs and motor outputs, emotions, attention, explicit or working memory, self-reflection, and language are all capabilities that have no roles in consciousness or in robot consciousness.

Edlund et al. (2011) performed artificial life experiments to analyze the evolution of simple agents aimed to solve a maze in a simulated environment. The authors found a clear correlation between the measures of information integration and the measures of fitness of the agent, suggesting that information integration capabilities evolve and are related to the functional complexity of the agent.

23.5 Self-Consciousness in Robots and Machines

A significant topic of robot consciousness is to give a robot the capabilities of self-awareness—that is, to reflect about itself, its perceptions, and actions during its operating life. According to this approach, a computational model of the mind may be made up of a hierarchy of modules, where low-level modules are related to reactive input-outputs, and middle-level modules are related to deliberative planning and reasoning. The high-level modules are associated with self-monitor and self-reflection capabilities.

The first theoretically founded attempt to give self-reflection capabilities to an artificial reasoning system is described in the seminal paper of Weyhrauch (1980). Weyhrauch proposed the reasoning system FOL, able to perform inferences and based on a logic system and a simulation structure capable of analog representations. The system can exploit meta

representations and reflect about itself, its inferences, and its capabilities. Weyhrauch (1995) discusses the relationships between FOL and consciousness in artifacts. The original implementation of FOL is still available in LISP (see link in the list of additional resources).

An early attempt to model consciousness by considering different levels of representation is in Johnson-Laird (1983). In the well-known book on mental models, Johnson-Laird discusses consciousness as the “operating system” of the mind. Several unconscious distributed processes run in the brain, and consciousness acts as the central control system of the mind, a sort of operating system. According to this view, the content of consciousness is made up of the value parameters of the central control system.

Minsky (2006) described a multiagent system based on several interacting agents at different levels, in which the tasks of higher-level agents are self-reflection and self-consciousness (figure 23.3). In detail, Minsky proposed different levels of agents, in which each level reflects on and critiques the levels beneath.

The first levels of the system are related to agents devoted to instinctive reflexes and learned reactions. The middle level is relevant to deliberation—that is, to the prediction-planning capabilities of the system. The higher levels are related to reflection, self-reflection, and self-consciousness. In particular, the reflection level is related to the ability to criticize the deliberative techniques adopted in the previous level; the self-reflection level is associated with




Figure 23.3

An outline of the multiagent system proposed by Minsky (2006).

the ability to generate critiques of the deficiencies and the weaknesses in the knowledge and methods employed by the system.

The higher level of the system is related to self-consciousness—that is, the ability to reflect on what others may think of the capabilities and performances of the system itself. A first attempt to implement the scheme proposed by Minsky in a simulated world was described by Singh and Minsky (2005).

Sloman and Chrisley (2003) followed a similar approach in the design of the H-CogAff architecture. H-CogAff is a framework architecture based on three primary levels related to reactive mechanisms, deliberative reasoning, and metamanagement—that is, reflective processes. The proposed framework prescribes different types of information, forms of representation, uses of data and types of mechanism for each level, and ways to put them together in the architecture. The SimAgent Toolkit is a freely available implementation in the Poplog framework.

McDermott (2001) made a distinction between *normal* access to the output of a computational module and *introspective* access to the same module. The first concerns the output related to the processing algorithms of the module. The second is related to the higher-order access within the processing of the module according to the self-model. He discussed the relationships between higher-order access and phenomenology in the line of higher-order theories of consciousness (see, e.g., Carruthers 1996).

McCarthy (1995) stressed the idea that a robot needs the ability to observe its mental states. He proposed a logic formalism to deal with aspects of self-reflection that could make robots conscious of their mental states. In detail, he presented the “mental situation calculus,” an extension of the situation calculus formalism aimed at modeling introspective actions in robots.

According to the classic version of situation calculus (see, e.g., Reiter 2001), the evolution of a state of affairs in the world is modeled by a sequence of situations $S_0, S_1, S_2, \dots, S_n$. The world changes when an instantaneous action a is performed. A new situation S_i is the result of the application of action a to the old situation S_{i-1} ; then $S_i = \text{Result}(a, S_{i-1})$. In the situation calculus formalism, the truth value of a proposition p depends on the considered situation. Then the formula $\text{Holds}(p, S_i)$ means that p is true in the situation S_i .

Let us consider the situation S_i where the robot knows the proposition p —for example, the color of the object A. The formula $\text{Holds}(\text{Know}(\text{Color}(A)), S_i)$ formalizes the fact that the robot knows the color of A. The situation in which the robot infers by introspection that it does not know the color of A is formalized by the formula $\text{Holds}(\text{Know}(\text{Not}(\text{Know}(\text{Color}(A)))), S_i)$. In this case, the robot knows that it does not know the color of A. Then, because of this fact, the robot may start some actions to learn the color of A.

The mental state of the robot may evolve because of learning actions. Let us consider the previous mental situation S_i in which the robot does not know the color of A. As an effect of teaching activities, the robot may learn the color of A. Then its mental state evolves to a new situation: $S_{i+1} = \text{Result}(\text{Learn}(\text{Color}(A)), S_i)$. The robot is in a new mental situation in which it now knows the color of A: $\text{Holds}(\text{Knows}(\text{Color}(A)), \text{Result}(\text{Learn}(\text{Color}(A)), S_i))$. Forgetting actions may be modeled similarly.

The mental situation calculus wants to capture the dynamics of self-reflection so that a robot may reason about its mental states. As emerges from the previous examples, the propositions and actions are mental, and the situations are the mental states of the robot.

In summary, the mental situation calculus is aimed at capturing the dynamic evolution of robot mental states.

Chella et al. (2008) proposed a cognitive architecture for a robot with introspective capabilities, organized in three computational areas. The *subconceptual* area is concerned with the low-level processing of perceptual data coming from the sensors. In the *linguistic* area, representation and processing are based on a logic formalism. In the conceptual area, the data coming from the subconceptual area are organized in *conceptual* categories.

Robot self-consciousness is based on the higher-order perception of the robot, in the sense that the first-order perception of the robot is the immediate perception of the environment, while higher-order perception is the perception of the inner world of the robot.

The described cognitive architecture has been tested on the board of a moving robot performing guided tours at the Archaeological Museum of Agrigento, Italy.

23.6 Global Workspace Theory

The global workspace theory (GWT) was proposed by Baars (see, e.g., Baars 1997) as the unification of different processes in the cortex. The GWT is tightly related to the global neuronal theory discussed by Dehaene et al. (2003). Baars observed that the brain could perform an enormous amount of unconscious parallel processing, while consciousness is serial and of limited capacity.

The GWT is based on assumptions that the brain is a collection of many specialized processors. Consciousness is associated with a *global workspace* whose contents “broadcast” to the processors. The processors work in parallel, and they compete to gain access to the global workspace (figure 23.4, *left*).

At some point, one processor wins the competition, and it gains access to the global workspace. Then it enters into consciousness and broadcasts to all the other processors to recruit others and to select the corresponding action (figure 23.4, *right*).




Figure 23.4

Global workspace theory. *Left*: Several unconscious processors compete to gain access to the global workspace. *Right*: The winning processor gains access to the global workspace—that is, to consciousness—and it recruits other processors.

Let us consider, for example, an agent attending an elaborate scene where there are many moving objects. According to the GWT, every moving object may be processed by an unconscious processor. All processors compete to gain access to the global workspace. Then, at some point, one processor corresponding, for example, to a ball moving toward the agent wins the competition, and it enters into consciousness. The winning processor recruits other processors to select the best action to be performed: for example, it will recruit the processors related to the motion of the arm so that the arm catches the moving ball.

Contexts shape conscious contents, and they constrain the competition of unconscious processors. Therefore, a coalition of processors may be expedited to gain access in a particular context and to recruit other processors. For example, a context related to a specific emotion may assist processors in achieving consciousness instead of other processors.

The GWT is a framework theory, and several cognitive architectures inspired by the GWT have been proposed in the literature. The main cognitive architecture is LIDA (Learning Intelligent Distributed Agent), developed by Stan Franklin and colleagues over the years (see, e.g., Franklin et al. 2014; see also chapter 10 for a general discussion of cognitive architectures).

Baars and Franklin (2009) reported on the relationships between LIDA and the GWT. An initial version of LIDA, named IDA, was built by Franklin (2003) as a dispatching system for the US Navy. The goal of IDA was to assign sailors to new billets at the end of their tours of duty. These assignments were performed by detailers, and IDA completely automated the roles of detailers. Interaction with sailors was performed by email in natural language, and IDA was able to negotiate the new billets with sailors and to write orders to them.

An overview of LIDA is shown in figure 23.5. Several processors based on different technologies were implemented in the architecture, such as neural networks, sparse distributed memories, schema mechanisms, behavior networks, and subsumption architectures. LIDA performs several aspects of the GWT, like perception, attention, episodic and declarative memories, the global workspace, and the selection of actions.

The cognitive cycle of LIDA is based on the following steps:

- The system perceives an entity, giving rise to a percept.
- The percept is sent to a preconscious buffer, where the percept gives rise to local associations.
- The percept competes for consciousness.
- If the percept wins the competition, then it broadcasts to all the other processors to recruit for resources.
- An action is selected according to the goal context hierarchy.
- Once the action is selected, then the action is executed, and the cognitive cycle restarts.

The chosen action may be performed immediately, or it may be sent back to the perceptual system for further examinations.

The LIDA architecture presents learning capabilities through the feedback generated by the global workspace. The feedback signals are sent to the unconscious modules, and they provide the basis of the reinforcement- and associative-learning processes of the architecture. The Lidapy framework is a freely available recent implementation of LIDA in Python (see link in additional resource list).

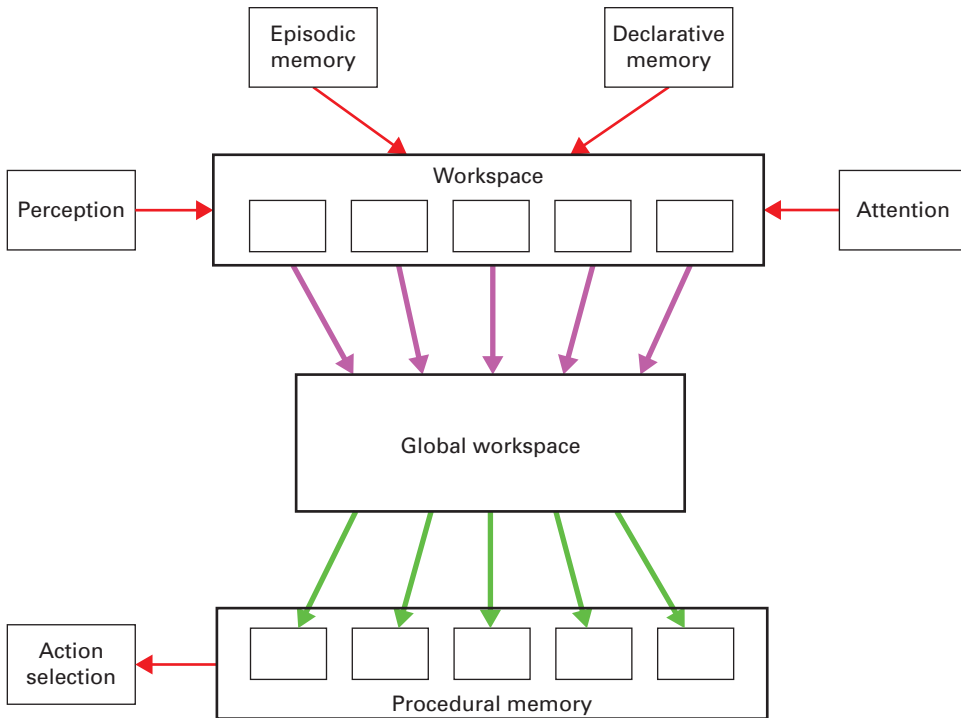


Figure 23.5
An overview of the LIDA cognitive architecture.

The LIDA architecture has proved to fit a body of empirical evidence concerning consciousness. Notably, a version of LIDA (Madl et al. 2011) implementing the Allport (1968) test modeling the phenomenal simultaneity of stimuli obtained time frames comparable to human subjects. Ramamurthy and Franklin (2009) discuss the general problems of conscious experiences and functional consciousness in the framework of LIDA.

Other cognitive architectures inspired by the GWT have been proposed in the literature. Shanahan (2006) discussed a cognitive architecture for a robot that extends the GWT by considering a cognitive cycle made up of an inner and an outer loop. The outer loop is similar to the cycle previously discussed in LIDA, while the role of the inner loop is to simulate the interaction with the environment internally. The internal simulation facilitates anticipation and planning in the architecture: the robot may internally simulate the effects of the actions before choosing the current course of activities.

Arrabales et al. (2009) discussed CERA-CRANIUM, a cognitive architecture based on GWT that controls a video game character. The architecture performed well in the BotPrize competition (Hingston 2009), a kind of Turing test (see below) in which autonomous bots have to convince a jury that they are human controlled. Notably, the CERA-CRANIUM bot won the award for the most humanlike bot at the 2010 competition. The software code of the bot is freely available (see link in the additional resource list).

Haikonen (see, e.g., Haikonen 2019), starting from engineering principles, designed the HCA, or Haikonen cognitive architecture, which presents contact points with the GWT.

The HCA is at the basis of the operating robot XCR-1, where many modules are implemented, including the auditory module, the visual module, and the emotional module. The modules send broadcast signals and compete in a winner-takes-all fashion to control the robot, similar to GWT. XCR-1 presents many aspects of machine consciousness: the robot can self-talk, respond to visual stimuli, and “feel” pain and emotions, among other functionalities.

23.7 The Internal Model Hypothesis

The internal model hypothesis states that an agent, to act in an intelligent and meaningful way, operates via an internal model of itself and the external world. The internal model allows the agent the capability to simulate its actions and evaluate its outcomes before doing them in the external environment. In this way, the agent can generate expectations about the course of events in the world and on the outcomes of its actions.

The internal model hypothesis is inspired by the “small-scale model” of reality discussed by Craik (1943). Dennett (1996) discusses “Popperian” creatures—that is, creatures able to generate theories about the external world and simulate experiments in their internal environment.

The proposal of an internal model acting as a simulation structure in a robot is not new: robot architectures have been proposed in the literature that present forms of an internal model of themselves and the external environment. Early examples have been provided by Mel (1990), Stein (1994), and Payton (1990), among others.

According to Hesslow (2002), the internal model hypothesis allows the brain to simulate actions, to simulate perceptions, and to generate anticipation about future events. Hesslow claims that conscious thoughts are based on these simulations. As the simulation of perception is related to the internally generated sensory inputs resembling the perception of the external world, it would be accompanied by the experience of the internal model of the world.

In brief, the internal model hypothesis states that consciousness arises from interaction between the internal model of the agent and the internal model of the world. Let us consider an agent interacting with the external world (figure 23.6, *top*).

Let us now consider the internal model of the agent, including the model of the agent and the external world (figure 23.6, *bottom*). According to the internal model hypothesis, consciousness arises not from the interaction of the agent with the external world but instead from the interaction of the internal model of the agent with the internal model of the external world. Susan Blackmore (1986) states that “being conscious is simply what it is like being a representation of the world” (163).

Figure 23.7, inspired by Grush (2004) and Gerdes and Happee (1994), describes the general framework of the internal model. A similar structure has been presented by Gray (2006). The robot has an internal model of itself and the external environment, allowing it to simulate its interactions with the external world. The controller sends the control signal at the same time to the real robot moving in the external world and to the inner model of the robot moving in the inner environment. Again, according to the internal model hypothesis, robot consciousness arises in the interaction of the internal model of the robot with the internal model of the situation.




Figure 23.6
The internal model hypothesis. *Top:* The agent interacting with the external world. *Bottom:* The agent with an internal model of itself interacting with an internal model of the external world.

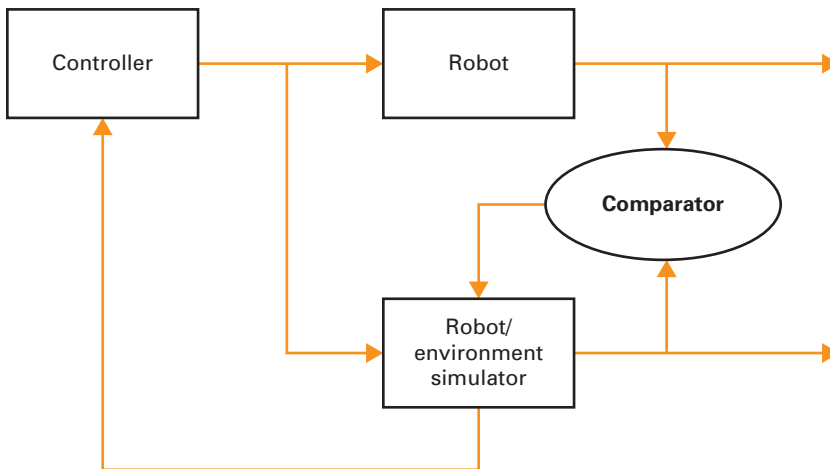


Figure 23.7
A general framework of the internal model hypothesis for robot consciousness.

A robot implementation inspired by the internal model hypothesis is *EcceRobot*, developed by Holland and colleagues (Holland 2007; Holland et al. 2007). *EcceRobot* is an anthropometric robot with a humanlike body. The robot has an internal simulator of itself and the environment that is able to represent in three dimensions (3D) the robot and the environment. The internal 3D simulation is employed to teach suitable neural networks how to control the motors of the robot.

Bongard et al. (2006) describe a “starfish” robot, a four-legged robot that generates a 3D model of itself by trial and error using suitable genetic algorithms. The robot uses the actuation-sensation relationship to infer an internal model of its body, and then it uses this model to learn locomotion. The robot is resilient: in case of damage—for example, a broken leg—the robot can generate a new model of its body and learn locomotion again with its current damaged body. A similar approach was described by Cully et al. (2015).

Chella and Macaluso (2009) discussed the robot *CiceRobot*, which was able to offer guided tours in an indoor and outdoor museum and was based on the internal model hypothesis. The architecture was instantiated on a wheeled robot for indoor and outdoor use. Currently, it is instantiated on a Pepper robot. The robot is a case study of many capabilities associated with the functional aspects of consciousness: to build and to maintain an internal model of the environment and itself, to pay attention to the relevant entities in the environment, to integrate information from different sources and different parts of the same source, to generate expectations about the possible events in the environment, to self-monitor, to simulate emotional states, and to process information by making it globally available to the robot.

The primary outcome of the case study was the acceptancy and transparency of the autonomous behavior of the robot in an environment populated by untrained users as museum tourists.

23.8 Tests for Robot Consciousness

People are concerned that current robot systems might already be conscious, so a substantial amount of research has been conducted on how a robot system can be tested for consciousness. An extended review of proposed criteria for consciousness in machines and robots is discussed by Elamrani and Yampolskiy (2019).

Many tests are based on the famous Turing (1950) test of imitation, in which a human interrogates an entity by teletype and decides whether they are examining a human or a machine that imitates human responses.

Sloman (2010) proposed the Robot Philosopher Test, a variant of the Turing test in which the arguments of discussion between the human tester and the tested entity are the philosophical theories of consciousness and experience.

Schneider and Turner (2017; see also Turner and Schneider 2019) proposed the Artificial Consciousness test (ACT), another variant of the Turing test in which the questions to be posed are focused on the quality of the inner experience of the entity under examination. The entity must be isolated from the external world to avoid the risk that a smart machine may retrieve the correct answers from the internet.

Harnad (1991) extended the Turing test by proposing the Total Turing test, in which a robot—that is, an embodied entity—should imitate the whole of human behavior in different situations.

Another source of inspiration for consciousness tests is the mirror test for primates (Gallup 1970; Gallup et al. 2002). In this case, a robot should recognize and describe itself and its movements by looking through a mirror, even in the presence of other robots and distractors. See Gold and Scassellati (2005), Chella et al. (2003), Suzuki et al. (2005), and Haikonen (2007) for examples of robot implementations of the mirror test.

Consciousness in robots and machines can be assessed by measuring specific features ascribed to consciousness, like the ability to presents forms of creativity—that is, to produce something new and unexpected. Bringsjord et al. (2001) presented the Lovelace Test, named after Ada Lovelace, while Chella and Manzotti (2012) discussed how a conscious robot should be able to improvise jazz in a jazz ensemble.

A related approach is to consider the capability of the conscious robot to generate a genuine inner speech, as proposed by Haikonen (2007). Inner speech is considered tightly related to self-consciousness (Morin 2005). Steels (2003), Clowes (2007), Arrabales (2012), and Chella et al. (2020) demonstrate examples of robots presenting forms of inner speech.

Another approach for testing machine consciousness is to apply the algorithmic theories proposed for human and robot consciousness, such as the previously described set of axioms by Aleksander and Dunmall or the $\Phi(S)$ measure derived from the information integration theory.

Gamez (2010) implemented SpikeStream, a freely available neural network simulator able to measure the $\Phi(S)$ of different kinds of networks (see the link to the system in the list of additional resources). In detail, Gamez applied $\Phi(S)$ to analyze the neural networks at the basis of SIMNOS, a simulation of EcceRobot.

Iklé et al. (2019) followed a similar approach to measure $\Phi(S)$ in the cognitive system controlling the robot Sophia when the robot was reading and when it was conversing. Seth et al. (2006) and Gamez and Aleksander (2009) proposed methods for designing suitable neural networks presenting high values of the measure $\Phi(S)$.

An interesting approach to assess consciousness in robots and machines was proposed by Arrabales et al. (2010a). They discussed ConsScale, a scale of consciousness in artificial agents that scores from -1 and 0 (the disembodied and isolated agent) to 11 (the super-conscious agent).

ConsScale considers a generic characterization of an artificial agent to comprise a body, a set of sensors, a set of actuators, a set of software routines, types of memories, and an external environment where the agent operates.

ConsScale assigns a level of consciousness according to the architectural complexity of the agent and to the behaviors of the agent. At the low level of ConsScale are reactive agents based on a direct link between sensors and actuators. At the intermediate levels are the agents able to adapt themselves, to pay attention, to generate plans, and to have emotions. At the higher level of the scale are the self-conscious agents, the empathic agents, and the social agents. At the top level is the humanlike agent, which can pass the Turing test, and the superconscious agent, able to manage several streams of consciousness. The ConsScale calculator is freely available (see the link in the list of additional resources).

Arrabales et al. (2010b) tested ConsScale by assessing some cognitive architectures such as CERA-CRANIUM, CRONOS (an implementation of EcceRobot), LIDA, and a version of the HCA. According to the assessment by ConsScale, the HCA and LIDA

received the highest score because they were successful at the emotional level—that is, at an intermediate level of consciousness. No architectures entered the higher levels.

23.9 Conclusion

Chella and Manzotti (2009) wrote a manifesto for robot consciousness in which they discussed some of the main challenges in the field. Notwithstanding the progress in this field, as seen in the numerous machine consciousness theories presented above, the challenges from this manifesto are still valid today. They include the role of embodiment and situatedness in machine consciousness, the roles of emotion and motivation, the difficulties in achieving information integration, the concept of time for robot consciousness, the question of free will for robots, and finally, the issue of robot experience.

The possible advent of conscious robots would lead to ethical concerns as well as issues related to the social integration of such robots. Bryson (2012, 2018) discussed in detail the risks of our moral obligations toward self-conscious systems. According to Bryson (2018, 15), “While constructing AI systems as either moral agents or patients is possible, neither is desirable.”

According to Gunkel (2012), if an entity has subjective experiences and is capable of suffering, then it should be treated as a person. These arguments may force us to review our fundamental definitions of the concept of person. If we assert that a robot system is conscious, then the moral responsibility of the system for its actions must be recognized. On the other hand, we may have to concede moral rights to conscious robots, such as the right to not be switched off.

In summary, robot consciousness is a research field that not only offers outstanding opportunities but brings ethical risks that cannot be undervalued.

Additional Reading and Resources

- This collection of classic papers on machine and robot consciousness is a valuable academic reference in the field: Chella, A., and R. Manzotti, eds. 2007b. *Artificial Consciousness*. Exeter, UK: Imprint Academic.
- This book is an introduction to robot consciousness from the perspectives of philosophy, cognitive science, and computer science, written by a founding father of the discipline: Aleksander, I. 2015. *Impossible Minds: My Neurons, My Consciousness*. Rev. ed. Singapore: World Scientific.
- This freely available e-book is a collection of papers that cover the most recent research trends of consciousness in robots and AI systems: <https://www.frontiersin.org/research-topics/5781/consciousness-in-humanoid-robots>. Chella, A., A. Cangelosi, G. Metta, and S. Bringsjord, eds. 2019. *Consciousness in Humanoid Robots*. Lausanne: Frontiers Media. doi:10.3389/978-2-88945-866-0.
- This new journal, with a freely available inaugural issue, presents the latest works in the field of consciousness in robotics and AI: <https://www.worldscientific.com/worldscinet/jaic>.
- HyperSlate™ logical framework by Bringsjord: <https://rpi.logicmodernapproach.com/>.

- Reasoning system FOL by Weyhrauch: <https://github.com/getfol/GETFOL>.
- The SimAgent Toolkit by Aaron Sloman: <https://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html>.
- The LIDA framework: <https://github.com/CognitiveComputingResearchGroup/lidapy-framework>.
- The CERA-CRANIUM bot: <https://github.com/raul-arrabales/CCbot4>.
- The SpikeStream simulator by Gamez: <http://spikestream.sourceforge.net/>.
- The ConsScale consciousness calculator by Arrabales: https://www.conscious-robots.com/conscale/calc_30.html.

Note

1. Swartz Foundation, *Final Report of the Workshop Can a Machine Be Conscious*, 2001, http://www.theswartzfoundation.org/abstracts/2001_summary.asp.

References

- Aleksander, Igor. 1992. "Capturing Consciousness in Neural Systems." In *Artificial Neural Networks 2, Proceedings of ICANN 1992 Conference*, 17–22. Amsterdam: Elsevier.
- Aleksander, Igor. 2005. *The World in My Mind, My Mind in the World*. Exeter, UK: Imprint Academic.
- Aleksander, Igor. 2015. *Impossible Minds: My Neurons, My Consciousness*. Revised ed. Singapore: World Scientific.
- Aleksander, Igor, and Barry Dunmall. 2003. "Axioms and Tests for the Presence of Minimal Consciousness in Agents." *Journal of Consciousness Studies* 10 (4–5): 7–18.
- Allport, David A. 1968. "Phenomenal Simultaneity and the Perceptual Moment Hypothesis." *British Journal of Psychology* 59:395–406.
- Arrabales, Raúl. 2012. "Inner Speech Generation in a Video Game Non-player Character: From Explanation to Self?" *International Journal of Machine Consciousness* 4 (2): 367–381.
- Arrabales, Raúl, Agapito Ledezma, and Araceli Sanchis. 2009. "Towards Conscious-Like Behavior in Computer Game Characters." In *Proceedings of the IEEE International Conference on Computational Intelligence and Games*, 217–224. Piscataway, NJ: IEEE Press.
- Arrabales, Raúl, Agapito Ledezma, and Araceli Sanchis. 2010a. "ConsScale: A Pragmatic Scale for Measuring the Level of Consciousness in Artificial Agents." *Journal of Consciousness Studies* 17 (3–4): 131–164.
- Arrabales, Raúl, Agapito Ledezma, and Araceli Sanchis. 2010b. "The Cognitive Development of Machine Consciousness Implementations." *International Journal of Machine Consciousness* 2 (2): 213–225.
- Baars, Bernard J. 1997. *In the Theater of Consciousness: The Workspace of the Mind*. Oxford: Oxford University Press.
- Baars, Bernard J., and Stan Franklin. 2009. "Consciousness Is Computational: The LIDA Model of Global Workspace Theory." *International Journal of Machine Consciousness* 1 (1): 23–32.
- Balduzzi, David, and Giulio Tononi. 2008. "Integrated Information in Discrete Dynamical Systems: Motivation and Theoretical Framework." *PLoS Computational Biology* 4 (6): e1000091. <https://doi.org/10.1371/journal.pcbi.1000091>
- Blackmore, Susan. 1986. "What It's Like to Be a Mental Model." In *Research in Parapsychology*, edited by D. Weiner and D. Radin, 163–164. Metuchen, NJ: Scarecrow.
- Blackmore, Susan, and Emily T. Troscianko. 2018. *Consciousness—an Introduction*. London: Routledge.
- Bongard, Josh, Victor Zykov, and Hod Lipson. 2006. "Resilient Machines through Continuous Self-Modeling." *Science* 314:1118–1123.
- Bringsjord, Selmer. 2007. "Offer: One Billion Dollars for a Conscious Robot. If You're Honest, You Must Decline." *Journal of Consciousness Studies* 14 (7): 28–43.
- Bringsjord, Selmer, Paul Bello, and David Ferrucci. 2001. "Creativity, the Turing Test, and the (Better) Lovelace Test." *Minds and Machines* 11:3–27.

- Bringsjord, Selmer, Paul Bello, and Naveen Sundar Govindarajulu. 2018. "Toward Axiomatizing Consciousness." In *The Bloomsbury Companion to the Philosophy of Consciousness*, edited by D. Jacquette, 289–324. London: Bloomsbury Academic.
- Bringsjord, Selmer, John Licato, Naveen Sundar Govindarajulu, Rikhiya Ghosh, and Atriya Sen. 2015. "Real Robots That Pass Human Tests of Self-Consciousness." In *24th IEEE International Symposium on Robot and Human Interactive Communication*, 498–504. Piscataway, NJ: IEEE Press.
- Bryson, Joanna. 2012. "A Role for Consciousness in Action Selection." *International Journal of Machine Consciousness* 4 (2): 471–482.
- Bryson, Joanna. 2018. "Patience Is Not a Virtue: The Design of Intelligent Systems and Systems of Ethics." *Ethics and Information Technology* 20:15–26.
- Carruthers, Peter. 1996. *Language, Thought and Consciousness: An Essay in Philosophical Psychology*. Cambridge: Cambridge University Press.
- Chalmers, David J. 1995. "Facing Up to the Problem of Consciousness." *Journal of Consciousness Studies* 2 (3): 200–219.
- Chella, Antonio, Angelo Cangelosi, Giorgio Metta, and Selmer Bringsjord. 2019. "Editorial: Consciousness in Humanoid Robots." *Frontiers in Robotics and AI* 6:17. <https://doi.org/10.3389/frobt.2019.00017>.
- Chella, Antonio, Marcello Frixione, and Salvatore Gaglio. 2003. "Anchoring Symbols to Conceptual Spaces: The Case of Dynamic Scenarios." *Robotics and Autonomous Systems* 43:175–188.
- Chella, Antonio, Marcello Frixione, and Salvatore Gaglio. 2008. "A Cognitive Architecture for Robot Self-Consciousness." *Artificial Intelligence in Medicine* 44:147–154.
- Chella, Antonio, and Irene Macaluso. 2009. "The Perception Loop in Cicerobot, a Museum Guide Robot." *Neurocomputing* 72:760–766.
- Chella, Antonio, and Riccardo Manzotti, eds. 2007a. *AI and Consciousness: Theoretical Foundations and Current Approaches, Papers from the 2007 AAAI Fall Symposium*. Menlo Park, CA: AAAI Press.
- Chella, Antonio, and Riccardo Manzotti, eds. 2007b. *Artificial Consciousness*. Exeter, UK: Imprint Academic.
- Chella, Antonio, and Riccardo Manzotti. 2009. "Machine Consciousness: A Manifesto for Robotics." *International Journal of Machine Consciousness* 1 (1): 33–51.
- Chella, Antonio, and Riccardo Manzotti. 2012. "Jazz and Machine Consciousness: Towards a New Turing Test." In *Revisiting Turing and His Test: Comprehensiveness, Qualia, and the Real World*, edited by Vincent C. Müller and Aladdin Ayesh, 49–53. Birmingham, UK: AISB/IACAP.
- Chella, Antonio, Arianna Pipitone, Alain Morin, and Famira Racy. 2020. "Developing Self-Awareness in Robots via Inner Speech." *Frontiers in Robotics and AI* 7:16. <https://doi.org/10.3389/frobt.2020.00016>
- Clowes, Robert. 2007. "A Self-Regulation Model of Inner Speech and Its Role in the Organisation of Human Conscious Experience." *Journal of Consciousness Studies* 14 (7): 59–71.
- Clowes, Robert, Steve Torrance, and Ron Chrisley. 2007. "Machine Consciousness: Embodiment and Imagination." *Journal of Consciousness Studies* 14 (7): 7–14.
- Craik, Kenneth J. W. 1943. *The Nature of Explanation*. Cambridge: Cambridge University Press.
- Cully, Antoine, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. "Robots That Can Adapt Like Animals." *Nature* 521:503–507.
- Dehaene, Stanislas, Hakwan Lau, and Sid Kouider. 2017. "What Is Consciousness, and Could Machines Have It?" *Science* 358:486–492.
- Dehaene, Stanislas, Claire Sergent, and Jean-Pierre Changeux. 2003. "A Neuronal Network Model Linking Subjective Reports and Objective Physiological Data during Conscious Perception." *Proceedings of the National Academy of Sciences USA* 100 (14): 8520–8525.
- Dennett, Daniel. 1996. *Darwin's Dangerous Idea*. New York: Simon and Schuster.
- Edlund, Jeffrey A., Nicolas Chaumont, Arend Hintze, Christof Koch, Giulio Tononi, and Christoph Adami. 2011. "Integrated Information Increases with Fitness in the Evolution of Animals." *PLoS Computational Biology* 7 (10): e1002236.
- Edelman, Gerald M., George N. Reeke, W. Einar Gall, Giulio Tononi, Douglas Williams, and Olaf Sporns. 1992. "Synthetic Neural Modeling Applied to a Real-World Artifact." *Proceedings of the National Academy of Sciences USA* 89:7267–7271.
- Elamrani, Aida, and Roman V. Yampolskiy. 2019. "Reviewing Tests for Machine Consciousness." *Journal of Consciousness Studies* 26 (5–6): 35–64.
- Floridi, Luciano. 2005. "Consciousness, Agents and the Knowledge Game." *Mind and Machines* 15:415–444.
- Franklin, Stan. 2003. "IDA—a Conscious Artifact?" *Journal of Consciousness Studies* 10 (4–5): 47–66.

- Franklin, Stan, Tamas Madl, Sidney D’Mello, and Javier Snaider. 2014. “LIDA: A Systems-Level Architecture for Cognition, Emotion, and Learning.” *IEEE Transactions on Autonomous Mental Development* 6 (1): 19–41.
- Gallup Jr., Gordon G. 1970. “Chimpanzees: Self-Recognition.” *Science* 167 (3914): 86–87.
- Gallup Jr., Gordon G., James R. Anderson, and Daniel J. Shillito. 2002. “The Mirror Test.” In *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, edited by M. Bekoff, C. Allen, and G. Burghardt, 325–333. Cambridge, MA: MIT Press.
- Gamez, David. 2010. “Information Integration Based Predictions about the Conscious States of a Spiking Neural Network.” *Consciousness and Cognition* 19 (1): 294–310.
- Gamez, David, and Igor Aleksander. 2009. “Taking a Mental Stance towards Artificial Systems.” In *Biologically Inspired Cognitive Architectures: Papers from the 2009 AAI Fall Symposium*. Menlo Park, CA: AAAI Press.
- Gamez, David, Zafeirios Fountas, and Andreas K. Fidjeland. 2013. “A Neurally-Controlled Computer Game Avatar with Human-Like Behavior.” *IEEE Transactions on Computational Intelligence and AI in Games* 5 (1): 1–14.
- Gerdes, V. G. J., and Riender Happee. 1994. “The Use of an Internal Representation in Fast Goal-Directed Movements: A Modeling Approach.” *Biological Cybernetics* 70:513–524.
- Gold, Kevin, and Brian Scassellati. 2005. “Learning about the Self and Others through Contingency.” In *Developmental Robotics: Papers from the 2005 AAI Spring Symposium*. Menlo Park, CA: AAAI Press.
- Gray, Jeffrey A. 2006. *Consciousness: Creeping Up on the Hard Problem*. Oxford: Oxford University Press.
- Grush, Rick. 2004. “The Emulator Theory of Representation: Motor Control, Imagery and Perception.” *Behavioral and Brain Sciences* 27:377–442.
- Gunkel, David J. 2012. *The Machine Question*. Cambridge, MA: MIT Press.
- Haikonen, Pentti O. 2007a. “Reflections of Consciousness: The Mirror Test.” In *AI and Consciousness: Theoretical Foundations and Current Approaches: Papers from the 2007 AAI Fall Symposium*, 67–71. Menlo Park, CA: AAAI Press.
- Haikonen, Pentti O. 2007b. *Robot Brains: Circuits and Systems for Conscious Machines*. Hoboken, NJ: John Wiley and Sons.
- Haikonen, Pentti O. 2019. *Consciousness and Robot Sentience*. 2nd ed. Singapore: World Scientific Press.
- Harnad, Stevan. 1991. “Other Bodies, Other Minds: A Machine Incarnation of an Old Philosophical Problem.” *Minds and Machines* 1 (1): 43–54.
- Hesslow, Germund. 2002. “Conscious Thought as Simulation of Behavior and Perception.” *Trends in Cognitive Sciences* 6 (6): 242–247.
- Hingston, Philip. 2009. “The 2K BotPrize.” In *Proceedings of IEEE International Conference on Computational Intelligence and Games*, 1–1. Piscataway, NJ: IEEE Press.
- Holland, Owen. 2003a. *Machine Consciousness*. Exeter, UK: Imprint Academic.
- Holland, Owen. 2003b. “Robots with Internal Models—a Route to Machine Consciousness?” *Journal of Consciousness Studies* 10 (4–5): 77–109.
- Holland, Owen. 2007. “A Strongly Embodied Approach to Machine Consciousness.” *Journal of Consciousness Studies* 14 (7): 97–110.
- Holland, Owen, Rob Knight, and Richard Newcombe. 2007. “A Robot-Based Approach to Machine Consciousness.” In *Artificial Consciousness*, edited by A. Chella and R. Manzotti. Exeter, UK: Imprint Academic.
- Iklé, Matthew, Ben Goertzel, Misgana Bayetta, George Sellman, Comfort Cover, Jennifer Allgeier, Robert Smith, et al. 2019. “Using Tonomi Phi to Measure Consciousness of a Cognitive System While Reading and Conversing.” In Vol. 2287, *Towards Conscious AI Systems: Papers of the AAI 2019 Spring Symposium*. Palo Alto, CA: CEUR Workshop Proceedings. <http://ceur-ws.org/vol-2287/paper20.pdf>.
- Johnson-Laird, Philip N. 1983. *Mental Models: Towards a Cognitive Science of Language, Inference and Consciousness*. Cambridge: Cambridge University Press.
- Koch, Christof. 2009. “A Theory of Consciousness.” *Scientific American Mind*, July/August, 16–19.
- Koch, Christof, Marcello Massimini, Melanie Boly, and Giulio Tononi. 2016. “Neural Correlates of Consciousness: Progress and Problems.” *Nature Reviews Neuroscience* 17 (5): 307–323.
- Koch, Christof, and Giulio Tononi. 2008. “Can Machines Be Conscious?” *IEEE Spectrum*, 47–51.
- Koch, Christof, and Giulio Tononi. 2017. “Can We Quantify Machine Consciousness?” *IEEE Spectrum*, 65–69.
- Krichmar, Jeffrey L., Douglas A. Nitz, Joseph A. Gally, and Gerald M. Edelman. 2005. “Characterizing Functional Hippocampal Pathways in a Brain-Based Device as It Solves a Spatial Memory Task.” *Proceedings of the National Academy of Sciences of the USA* 102 (6): 2111–2116.

- Madl, Tamas, Bernard J. Baars, and Stan Franklin. 2011. "The Timing of the Cognitive Cycle." *PLoS One* 6 (4): e14803. <https://doi.org/10.1371/journal.pone.0014803>.
- Mel, Bartlett. 1990. *Connectionist Robot Motion Planning*. Cambridge, MA: Academic Press.
- McCarthy, John. 1995. "Making Robots Conscious of Their Mental States." *Machine Intelligence* 15:3–17. <http://jmc.stanford.edu/articles/consciousness.html>.
- McDermott, Drew. 2001. *Mind and Mechanisms*. Cambridge, MA: MIT Press.
- Minsky, Marvin. 2006. *The Emotion Machine*. New York: Simon and Schuster.
- Morin, Alain. 2005. "Possible Links between Self-Awareness and Inner Speech." *Journal of Consciousness Studies* 12 (4–5): 115–134.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *Philosophical Review* 83 (4): 435–450.
- Oizumi, Masafumi, Larissa Albantakis, and Giulio Tononi. 2014. "From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0." *PLoS Computational Biology* 10 (5): e1003588. <https://doi.org/10.1371/journal.pcbi.1003588>.
- O'Regan, J. Kevin, and Alva Noë. 2001. "A Sensorimotor Account of Vision and Visual Consciousness." *Behavioral and Brain Sciences* 24:939–973.
- Payton, David W. 1990. "Internalized Plans: A Representation for Action Resources." *Robotics Autonomous Systems* 6 (1): 89–103.
- Ramamurthy, Uma, and Stan Franklin. 2009. "Resilient Architectures to Facilitate Both Functional Consciousness and Phenomenal Consciousness in Machines." *International Journal of Machine Consciousness* 1 (2): 243–253.
- Reeke, George N., Olaf Sporns, and Gerald M. Edelman. 1990. "Synthetic Neural Modeling: The 'Darwin' Series of Recognition Automata." *Proceedings of the IEEE* 78 (9): 1498–1530.
- Rees, Geraint, Gabriel Kreiman, and Christof Koch. 2002. "Neural Correlates of Consciousness in Humans." *Nature Reviews Neuroscience* 3 (4): 261–270.
- Reggia, James A. 2013. "The Rise of Machine Consciousness: Studying Consciousness with Computational Models." *Neural Networks* 44:112–131.
- Reggia, James A., Garrett E. Katz, and Gregory P. Davis. 2018. "Humanoid Cognitive Robots That Learn by Imitating: Implications for Consciousness Studies." *Frontiers in Robotics and AI* 5:1. <https://doi.org/10.3389/frobt.2018.00001>.
- Reiter, Raymond. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. Cambridge, MA: MIT Press.
- Rushby, John, and Daniel Sanchez. 2018. *Technology and Consciousness Workshop Report*. SRI International. <http://www.csl.sri.com/users/rushby/papers/techconscwks2017.pdf>.
- Schmidhuber, Juergen. 1992. "Learning Complex, Extended Sequences Using the Principle of History Compression." *Neural Computation* 4 (2): 234–242.
- Schneider, Susan, and Edwin Turner. 2017. "Is Anyone Home? A Way to Find Out if AI Has Become Self-Aware." *Scientific American* (blog). <https://blogs.scientificamerican.com/observations/is-anyone-home-a-way-to-find-out-if-ai-has-become-self-aware/>.
- Searle, John R. 2000. "Consciousness." *Annual Review of Neuroscience* 23:557–578.
- Seth, Anil K., Eugene Izhikevich, George N. Reeke, and Gerald M. Edelman. 2006. "Theories and Measures of Consciousness: An Extended Framework." *Proceedings of the National Academy of Sciences of the USA* 103 (28): 10799–10804.
- Shanahan, Murray P. 2006. "A Cognitive Architecture That Combines Internal Simulation with a Global Workspace." *Consciousness and Cognition* 15:433–449.
- Singh, Push, and Marvin Minsky. 2005. "An Architecture for Cognitive Diversity." In *Visions of Mind*, edited by D. Davis, 312–331. London: Idea Group.
- Sloman, Aaron. 2010. "Machine Consciousness: Response to Commentaries." *International Journal of Machine Consciousness* 2 (1): 75–116.
- Sloman, Aaron, and Ron Chrisley. 2003. "Virtual Machines and Consciousness." *Journal of Consciousness Studies* 10 (4–5): 133–172.
- Steels, Luc. 2003. "Language Re-entrance and the 'Inner Voice.'" *Journal of Consciousness Studies* 10 (4–5): 173–185.
- Stein, Lynn A. 1994. "Imagination and Situated Cognition." *Journal of Experimental and Theoretical Artificial Intelligence* 6 (4): 303–407.

- Suzuki, Tohru, Keita Inaba, and Junichi Takeno. 2005. "Conscious Robot That Distinguishes between Self and Others and Implements Imitation Behavior." In *International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems (IEA/AIE) 2005*, edited by M. Ali and F. Esposito, 101–110. LNAI 3533. Heidelberg: Springer.
- Tegmark, Max. 2016. "Improved Measures of Integrated Information." *PLoS Computational Biology* 12 (11): e1005123. <https://doi.org/10.1371/journal.pcbi.1005123>.
- Tononi, Giulio. 2004. "An Information Integration Theory of Consciousness." *BMC Neuroscience* 5:42. <https://doi.org/10.1186/1471-2202-5-42>.
- Tononi, Giulio. 2008. "Consciousness as Integrated Information: A Provisional Manifesto." *Biology Bulletin* 215:216–242.
- Tononi, Giulio, Melanie Boly, Marcello Massimini, and Christof Koch. 2016. "Integrated Information Theory: From Consciousness to Its Physical Substrate." *Nature Reviews Neuroscience* 17 (7): 450–461.
- Tononi, Giulio, and Christof Koch. 2008. "The Neural Correlates of Consciousness: An Update." *Annals of the New York Academy of Sciences* 1124:239–261.
- Tononi, Giulio, and Olaf Sporns. 2003. "Measuring Information Integration." *BMC Neuroscience* 4:31. <https://doi.org/10.1186/1471-2202-4-31>.
- Turing, Alan. 1950. "Computing Machinery and Intelligence." *Mind* 59 (236): 433–460.
- Turner, Edwin, and Susan Schneider. 2019. "Testing for Synthetic Consciousness: The ACT, the Chip Test, the Unintegrated Chip Test, and the Extended Chip Test." In Vol. 2287, *Towards Conscious AI Systems: Papers of the AAAI 2019 Spring Symposium*. Palo Alto, CA: CEUR Workshop Proceedings. <http://ceur-ws.org/vol-2287/short2.pdf>.
- Verschure, Paul. 2013. "From the Mirage of Intelligence to a Science and Engineering of Consciousness." *IEEE Intelligent Systems*, September/October, 7–10.
- Vimal, Ram L. P. 2009. "Meaning Attributed to the Term 'Consciousness'—an Overview." *Journal of Consciousness Studies* 16 (5): 9–27.
- Weyhrauch, Richard W. 1980. "Prolegomena to a Theory of Mechanized Formal Reasoning." *Artificial Intelligence* 13 (1–2): 133–170.
- Weyhrauch, Richard W. 1995. "Building Conscious Artifacts." In *Consciousness: Distinction and Reflection*, edited by G. Trautteur, 18–41. Napoli: Bibliopolis.
- Zylberberg, Ariel, Diego Fernández Slezak, Pieter R. Roelfsema, Stanislas Dehaene, and Mariano Sigman. 2010. "The Brain's Router: A Cortical Network Model of Serial Processing in the Primate Brain." *PLoS Computational Biology* 6 (4): e1000765. <https://doi.org/10.1371/journal.pcbi.1000765>.

Contributors

Yiannis Aloimonos, University of Maryland College Park, US

Minoru Asada, International Professional University of Technology in Osaka and Osaka University, Japan

Gianluca Baldassarre, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche, Italy

Michael Beetz, University of Bremen, Germany

Tony Belpaeme, Ghent University, Belgium, and University of Plymouth, UK

Angelo Cangelosi, University of Manchester, UK, and AIST-AIRC, Japan

Antonio Chella, Università degli Studi di Palermo and ICAR-CNR, Italy

Ravinder Dahiya, University of Glasgow, UK

Alessandro Di Nuovo, Sheffield Hallam University, UK

Marco Dorigo, IRIDIA, Université Libre de Bruxelles, Belgium

Diego Ferigo, Istituto Italiano di Tecnologia, Italy

Heiko Hamann, Institute of Computer Engineering, University of Lübeck, Germany

Mary Katherine Heinrich, IRIDIA, Université Libre de Bruxelles, Belgium

Guido Herrmann, University of Manchester, UK

Tiffany J. Hwu, HRL Laboratories LLC, US

Fumiya Iida, Cambridge University, UK

Yiming Jiang, Hunan University, China

Jeffrey L. Krichmar, University of California, Irvine, US

Ute Leonards, University of Bristol, UK

Erwin Jose Lopez Pulgarin, University of Manchester, UK

Vincent C. Müller, Eindhoven University of Technology, Netherlands, University of Leeds and Alan Turing Institute, UK

Shingo Murata, Keio University, Japan

Yukie Nagai, University of Tokyo, Japan

Lorenzo Natale, Istituto Italiano di Tecnologia, Italy

Stefano Nolfi, Institute of Cognitive Sciences and Technologies, National Research Council, Italy

Markellos Ntagios, University of Glasgow, UK

Tetsuya Ogata, Waseda University / AIST, Japan

Oliver Ozioko, University of Glasgow, UK

Alberto Parmiggiani, Istituto Italiano di Tecnologia, Italy

Jianxin Peng, Sichuan University, China

Daniele Pucci, Istituto Italiano di Tecnologia, Italy

Elena Rampone, Istituto Italiano di Tecnologia, Italy

Giulio Sandini, Istituto Italiano di Tecnologia, Italy

Kazuma Sasaki, DWANGO Co. Ltd., Japan

Luca Scimeca, Cambridge University, UK

Kuniyuki Takahashi, Preferred Networks Inc., Japan

Huajin Tang, Zhejiang University, China

Vadim Tikhanoﬀ, Istituto Italiano di Tecnologia, Italy

Silvio Traversaro, Istituto Italiano di Tecnologia, Italy

David Vernon, Carnegie Mellon University Africa, Rwanda

Mostafa Wahby, Institute of Computer Engineering, University of Lübeck, Germany, and IRIDIA, Université Libre de Bruxelles, Belgium

Jiru Wang, Sichuan University, China

Tatsuro Yamada, Panasonic Corp., Japan

Rui Yan, Zhejiang University of Technology, China

Chenguang Yang, Bristol Robotics Laboratory, University of the West of England, UK

Tom Ziemke, Linköping University, Sweden

Index

- AAAI Fall Symposium on Cognitive Robotics, 4, 12
- Abstract concepts, 401–402, 433–452
four domains, 437
neuroscience/psychology, 434–438
- Action representation, 403, 418, 428
- Action selection, 6, 20, 90, 193, 197, 257, 427, 438, 445, 464
- Actions structuring, 418–419
- ACT-R, 194, 196, 199, 207
- Adaptive control architecture, 341
- Affordance, 129, 131, 286, 401
- Aibo Robot, 50, 123–124, 128
- A.L.I.C.E. conversation agent, 403
- Altruistic Behavior, 135, 364–365, 370
- Android robots, 35, 37, 383–384, 403
- Animal-inspired soft robots, 100
- Ant colonies, 79
- Architectures. *See* Cognitive Architectures
- ARGoS simulator, 94
- Artificial consciousness. *See* Consciousness
- Artificial Consciousness Test (ACT), 467
- Artificial life, 9–10, 13, 459
- Asimov laws of robotics, 240
- Atari game, 71, 260, 262
- ATLAS robot, 136
- Attention, 88, 131, 191, 196, 197, 207, 257, 281–282, 287, 363–364, 441, 455. *See also* Joint attention
- Autism spectrum disorder (ASD), 222, 361, 371–372, 390
- Autobiographical memory, 204, 445
- Autoencoder, 69, 165, 176, 171–172, 179, 263, 316, 369, 405. *See also* Deep learning
- Automation and Employment, ethics, 234–236
- AutoMoDe, design, 85
- Autonomous Mental Development, 12, 32, 41
- Autonomous Systems, 237–239
- Autonomous vehicles, 237–238, 353–354
- Autonomous weapons, 238–239
- Babbling, language, 396–397
- Babbling, motor, 46, 70, 199, 261, 263
- BatSLAM, 302
- Baxter robot, 124–125, 127, 129, 130, 137, 321–329, 403
- BCI, manipulation, 316, 320–323
- Behavior-based robotics, 8–9, 10, 12, 84
- Biomimetic, 13, 22, 108, 145–159, 302
- Biomimetic skin (eskin), 145–159
- Blender simulator, 125
- Body Motion Query, 416–418
- Brain-Based Device, 23–24
- Braitenberg, vehicles, 9–10, 19–20
- CAN navigation, 301
- Care, robotics, 251, 233–234, 332, 339, 344
- CB2 robot, 11, 13
- CERA-CRANIUM cognitive architecture, 464, 468, 470
- Chatterbot, 403
- Choregraphe, robot software, 136
- CiceRobot robot, 467
- CLARION, 194, 196, 199, 207
- Coevolution, 66, 85, 223
- Cognition
attributes, 6
definition, 3
off-line, 7, 217
swarm, 87
theories, 218–221
- Cognitive abilities, core, 196–200
- Cognitive architectures, 123, 126, 191–213, 273, 289, 340–343, 349, 351, 354, 355, 444–445, 463–464
desirable characteristics, 195
consciousness, 456–457
- Cognitive consciousness, 456
- Cognitive Control, Decision, 337–356
- Cognitive control architecture, 342
- Cognitive dialogue, 282
- Cognitive map, Tolman, 295
- Cognitive map building, 300–306
- Cognitive science, foundations, 191–193
- Cognitive system, 3, 6, 13–14, 49, 191, 193–199, 217, 283, 468
- Cognitivist Paradigm, 192, 194
- COG robot, 8
- Collaboration, human-robot, 337–361. *See also* Human-robot interaction

- Collective
 - behavior, 65–67, 77–98
 - decision-making, 90–92
 - memory, 88
 - perception, 88
- Commonsense knowledge, 289, 413, 418, 424, 428
- Communication, 395–412. *See also* Language
 - direct, 82–83
 - indirect, 83–84
 - underwater, 84
- Competence-based intrinsic motivation, 50, 252–253, 257
- Complexity, 61, 84, 104–105, 112–113, 168, 182, 302, 338, 468
- Consciousness
 - machine, 453–474
 - neuroscience, 455
- ConsScale, consciousness scale, 468, 470
- Constructive developmental science, 374
- Control
 - cognitive, 337–360
 - robotics, 337–339
 - soft robot, 101–104
- Control in cognitive robotics and HRI, 343–344
- Conversational agents, 403
- Convolutional neural networks (CNNs), 165–166, 166–170, 286–287, 331, 405. *See also* Deep learning
- Counting gestures, 442
- CRAM Cognitive Architecture, 201–205, 207, 209, 429
- Crawling, 51
- Cyc ontology, 424
- Darwin robots, 20–26, 455
- Data collection, efficient, 181–183
- Deception, 233
- Decision and control action scheme (DCAS), 344–351, 353, 354
- Deep belief networks (DBNs), 405
- Deep learning (Deep neural networks), 166–181, 183, 208, 286–287, 427
 - language, 405–408
 - vision, 286–287
- Deep reinforcement learning, 179, 181, 208, 316, 330, 427
 - Grasping and Manipulation, 330–332
- Definition
 - abstractness, 433
 - cognition, 3, 196
 - cognitive robotics, 3–4, 12
 - consciousness, 453
 - embodiment, 215, 216
- Deontic cognitive event calculus (DCEC), 456–457
- Development, language, 395–397
- Development, nonlinear stages, 51, 52–53, 273, 395, 398, 438
- Developmental robotics, 6, 8, 12–13, 41–58, 99, 104, 107–111, 154, 197, 208, 252, 310, 374, 398–403, 442, 444
 - language, 398–403
 - soft robotics, 104–108
- Direction tuning, navigation, 298–299
- DolphinSLAM, 302
- Domain-adaptive meta-learning (DAML), 178
- Domain randomization, 181
- Driver assistance systems (Advanced DAS), 353
- Dynamical systems, 47–49, 107, 170, 274, 325
- Dynamic Movement Primitive (DMP), 318, 328–330, 332
- EASE project, 202–203, 207
- EcceRobot, 467, 468
- Education, abstract concepts, 434–438
- Education robotics, 389–390
- EEG/EMG, manipulation, 316, 318–320
- Electromagnetic actuators (EMAs), 155
- ELIZA, 403
- ELMER robot, 9, 19
- ELSIE robot, 9, 19
- Embodied cognition, 6–8, 12–13, 60, 193, 213–218, 224–227, 398, 402, 435, 438
- Embodiment, 5, 8, 14, 30, 36, 46–50, 105, 213–230, 272, 289, 396, 436–437, 441, 447, 469
 - abstract concepts, 433–436, 446
 - AI, 221–222
 - cognitive robotics, 222–225
 - cognitive science, 218–221
 - language, 396–397
 - learning through body, 113
 - vision, 272–274
- Emergent behavior, 47–48, 107–108, 112, 192
- Emergent paradigm, 192, 194
- eMODUL emotion system, 445
- Emotions, 206, 218, 368, 380, 387–389, 433–434, 438, 444–446, 453, 456, 465
- Empathy, artificial, 46, 52–49, 53
- Enactive cognition, 46–47, 49, 192, 222, 227
- Encyclopedic Knowledge Bases, 423–425
- Epigenetic Robotics (EpiRob), 12, 41, 52, 104, 252, 339
- Episodic memory, cognitive navigation, 308–309
- Epistemic intrinsic motivations (eIMs), 251
- ERA cognitive architecture, 399–400
- ERICA robot, 403
- Ethics, 42, 231–248, 453, 477
- Evolution, intrinsic motivation, 264–265. *See also* Intrinsic motivation
- Evolutionary algorithm, 60, 111–112
- Evolutionary robotics, 12–14, 22, 59–76, 85–89, 223
- Exploration and navigation, 309–310
- Extended Phenotype, 110
- External memory, 88
- Extrinsic motivations, 253–259
- FARSA simulator, 125
- Fault Tolerance, 81
- Fetus model, 5, 44–46, 52, 365
- Finger counting, 402–403, 442–444
- FOL reasoning system, 459
- Functional morphology, 105–107
- Gaussian Mixture Model (GMM), 173, 317–318, 332
- Gazebo simulator, 125–126, 135–136
- General control architecture, 338
- Generalized plan, knowledge base, 414, 418
- Global workspace theory (GWT), 455, 462–465
- Goal Formation, intrinsic motivation, 261
- Goal formation by imagination, 263
- Goal manifold search, 263

- Goal marking, 262
 Goal sampling, 261
 GOFAL, 3, 8, 219, 220, 221
 GRAIL cognitive architecture, 266
 Grasping, 69–71, 10, 104, 126, 130, 167–169, 203, 217, 255, 276, 278, 315–317, 330–333, 417
 Grid cell, 297–300, 306
 Grounding, language, 7, 14, 49, 170–172, 216, 218–220, 227, 285, 398–403, 407, 422, 436, 438–441, 446, 456
 Grounding transfer, 434, 438–441
 Group selection, 20, 92
 Growth, soft robot, 110
- Hand (in-hand) object manipulation, 170
 Hand (in-hand) object pose estimation, 169
 HCA cognitive architecture, 464–465, 468
 H-CogAff cognitive architecture, 461
 Head direction cell, 297, 299, 305
 Hebbian learning, 199, 263, 300–301, 399
 Hick's law, 93
 Hidden Markov Models (HMMs), 170–173, 386
 Historical Embodiment, 215
 History, Cognitive robotics, 10–14
 HOAP-3 robot, 178
 Honeybees, 92–93
 Human in the Loop, design, 84
 Humanoid robots, 8, 12, 20, 41, 112–113, 124–129, 134–138, 148, 205, 214, 216, 222–224, 233, 277, 289, 337, 351, 355, 382–384, 395, 408, 439. *See also* Robot platforms review
 Human-robot collaboration, 337–356. *See also* Human-Robot Interaction
 Human-Robot Interaction (HRI), 14, 123, 159, 207, 231, 233–234, 277–281, 332, 343–344, 379–394, 408
 applications, 389–390
 autonomous systems, 351–354
 control, 343–344
 dynamic Decision and Action Framework, 344–347
 ethics, 233–234
 neuroscience, 380–381
 non-verbal interaction, 386–388
 verbal interaction, 384–386
 vision, 277–281
 Hybrid Systems, 193, 195
 Hyperdimensional computing, 427–428
 HyperSlate logical framework, 468
- IBM Watson, 5, 425, 427
 iCub robot, 12, 69–71, 124, 126–129, 138, 278, 281, 369–370, 399–402, 439–440, 443–444, 445
 iCub simulator, 134–135
 Imitation, 46, 50, 89, 129, 137, 172–174, 176–179, 152, 222, 363, 368–369, 419, 467
 Imitation Learning, 172–179
 one-shot, 178–179
 Implicit/explicit social signs, 277
 Infant development, 46, 351
 Inference, for perception, 272
 Information integration theory (IIT) of consciousness, 457–458
 Inner speech, 453, 468
 Intelligent Robotics (AI robotics), 3
 Intention reading, 6, 362, 364–365, 369–370
- Interaction, robot-robot, 81–82, 88
 Internal model, 68–69, 344–345, 349–352, 317 hypothesis, 465–467
 Intrinsic motivation, 14, 46–47, 49–50, 179, 251–270
 Intrinsic tactile sensing, 148–152
 iRat robot, 125
 ISAC cognitive architecture, 197, 201, 205–207
 IsaacSim simulator, 135
- Joint attention, 46, 50–53, 279, 362–364, 367–368, 372, 374
- Kilobot, 78
 KISMET robot, 8, 444
 Knowledge-based, vision, 285
 Knowledge representation (knowledge based systems), 413–432
 KnowRob, 202–205, 425–427, 428
 Kuka LWR robot, 125–127, 129, 130, 136
- Language, 6, 49, 66, 125, 135, 165–166, 170–173, 200–201, 204–205, 213–215, 226, 277, 281–282, 285–287, 289, 311, 361, 384, 387, 395–412, 419–420, 433–447, 459
 Learning, 179–183. *See also* Reinforcement learning; Social learning
 from demonstration (LfD), 172, 328–329
 from play (Lp), 178–179
 cognitive capability, 197–198
 collective, 89–90
 imitation, 172–179
 One-Shot, 178–179
 Online, Open-Ended, Cumulative, 53
 social, 89–90
 Lexical, language analysis, 397–398
 LIDA cognitive architecture, 463–465, 468
 Lindenmayer systems, 80
 Logics, 420–421
 LSTM, 69, 166, 171, 176, 406–407. *See also* Recurrent neural networks
- Machine consciousness, 453–473. *See also* Consciousness
 Machine Ethics, 231, 239–242, 243
 Machine learning, 165–190. *See also* Deep learning
 Intrinsic motivation, 259–265
 robot language models, 404–407
 Majority rule, 91
 Manipulation, 6, 14, 48, 53, 63, 100, 113, 123–124, 127, 129, 168–169–171, 175, 179, 201, 315–336, 343, 351, 403, 413, 416
 EEG/EMG, 316, 318–320, 320–323
 in-hand object, 170
 MATLAB Robotics System Toolbox, 136
 Maturation, development, 51–52, 435
 Mental models, 105, 341, 460
 Metacognition, 196, 199–203, 205
 Micro-Macro Link, 78
 Microsoft Robotics Developer Studio, 132
 Minimal consciousness, 456
 Mirror neurons, 280, 363, 369, 435
 Mirror test, 468
 Model-agnostic meta-learning (MAML), 178
 Moral Agents, 240–241

- Morphological Computation, 7–8, 22, 49, 60–61, 106–107, 114
- Motion Learning, manipulation, 328–329
- MuJoCo simulator, 71, 132, 134, 137
- Multiagent system, Minsky, 460–461
- Multi-lateral latent Dirichlet allocation (MLDA), 404
- MultiNEAT, 94
- Multiple timescale recurrent neural networks (MTRNN), 166, 175–177, 401, 406
- Multirobot arm-picking, 181
- NAO robot, 123–124, 128–129, 136, 175–176, 382, 390, 406, 157
- Narrative-enabled episodic memories (NEEMS), 204
- Nature versus nurture, 51, 395–396
- Navigation, 6, 8, 20, 30–36, 82, 123, 126–127, 130, 200, 295–314, 343, 403
 - cognitive, 306–308
 - neuroscience/animal, 296–298
- Neural Darwinism, 20, 37
- Neuromorphic, 13, 23, 30–36, 459
- Neurorobotic Platform (Human Brain Project), 13, 136
- Neurorobotics, 9, 12–13, 19–40, 49, 85
- NeuroSLAM, 305, 311
- Neurosnake robot simulator, 125
- Neurosymbolic Learning and Reasoning, 427
- NLP, robot language, 404–405
- Nonverbal communication, 404
- Number learning, 49, 402–403, 441–444
- Object-action complexes (OAC), 419
- Object recognition, 23, 27, 168–169, 274–275, 285
- ODE Open Dynamics Engine simulator, 125–126, 133, 134–136
- Off-line cognition, 217
- Off-loading, cognitive, 64–65, 69
- One-Shot Learning, 174, 178–179
- Ontogenetic learning, 47, 51–52, 69–71, 99, 105–107, 110–111
- Ontologies, 203–204, 423–425
- OpenAI Gym, 183
- OpenCV, 289
- OpenEASE, 429
- Open-ended learning, 47, 53, 251–270, 368, 390, 408
- Organismic Embodiment, 216
- Organismoid Embodiment, 216
- ORO, knowledge representation, 425
- Outdoor robot swarm, 86–87
- Panda robot, 124, 127, 129–130, 137
- Path Integration, navigation, 298–300
- Perception, 5–6, 12, 37, 49, 51, 61, 88–89, 96, 108–109, 123–124, 128, 137, 146–147, 193, 196–197, 201–203, 207–208, 233–234, 271–289, 315, 340–342, 364, 387, 423, 425–426, 438, 456–457, 462. *See also* Sensing, tactile; Vision
- Phonetic, language analysis level, 397–398
- Phylogenetic-Ontogenetic Interaction, 51–52, 67–68
- Physical/Sensorimotor Embodiment, 216
- Piezoelectric/Capacitive Stack, 152–154
- Pioneer robot, 124–125, 127, 130
- Place cell, 300, 301
- Platforms, robot, 123–144
- Polar map, vision, 286, 288
- PolyScheme, cognitive architecture, 200
- Popperian creatures, 465
- PR2 robot, 178, 202, 284
- Pragmatic, language analysis level, 397–398
- Pragmatic everyday activity manifolds (PEAMs), 205
- PRAXICON cognitive architecture, 283, 403
- Predicate logic, 420–421
- Prediction, 6, 36, 50, 159, 174–175, 200, 252, 257–260, 275, 285, 369–372, 419, 460
- Prediction-based intrinsic motivation, 252
- Predictive Coding, 273, 370–371, 372
- Predictive Learning, 172–179
- Principles, Cognitive robotics, 6–9
- Principles, Developmental robotics, 46–53
- Privacy and Surveillance, ethics, 236–237
- Probabilistic representation and reasoning, 423
- Programming by demonstration (PbD), 172
- Prospection, 196, 200
 - vision, 274–276
- Prosthetic hand, 146, 152, 382
- Q-learning, 180, 199, 316, 331. *See also* Deep reinforcement learning
- QRIO robot, 174–176
- Question Answering, 282, 418–423
- RatSLAM, 30, 37, 301–302, 311
- Reaching, 48, 69–71, 131, 136, 175, 274, 278–280, 315–316, 330, 343, 355, 367, 515, 417
- Reactive cognition, 8, 63–64, 207, 273, 275
- Reasoning, 3–4, 8, 12, 191, 199, 201–205, 272–273, 340, 413–432, 433–434, 459–460, 470. *See also* Knowledge representation
- Recurrent neural networks (RNNs), 166, 171, 174–175, 356, 439, 442. *See also* Deep learning; LSTM
- REEM robot, 445
- Reinforcement learning, 71, 179–183. *See also* Deep Reinforcement Learning
 - joint attention, 367
- Representation, sparse, 36
- Representational-redescription model, 52
- Responsibility for robots, 241
- Rewards, Sparse, 259–260
- Rights for robots, 241–242
- RNNPB (Recurrent neural networks with parametric bias), 142, 174–176. *See also* Deep learning
- RoboCup, 8, 129
- Robot Philosopher Test, 467
- Robot platforms review, 127–131. *See also* Humanoid robots
- Robot-Robot Interaction, 81
- Robot simulator review, 131–137
- Robovie, 124–127, 129
- RobWorkSim simulator, 125
- Roomba robot, 8
- ROSETTA, knowledge representation, 425
- ROS middleware, 127, 134, 138, 356
- Salamander robot simulator, 125
- Sawyer robot, 178
- Scalability, 77, 80–81, 112–113, 121, 405
- SciBot robot, 382
- S-CTRNN Neural network, 371–372

- Segmentation, vision, 287
- Selection of Skills, intrinsic motivation, 263–264
- Self-Consciousness, 459–462
- Self-Organization, 13, 47–49, 77–80, 92, 216, 222
- Self-organizing map (SOM), 400, 404
- Self-Other Recognition, 362–363, 365–366
- Semantic, language analysis level, 397–398
- Sensing, tactile, 146–148, 167–170
- Sensing-Actuation integration, 154–158
- Sensorimotor coordination, 61–63
- Sensory EgoSphere (SES), 205
- Service robots, 214, 251, 310, 425
- Sex robotics, ethics, 234
- Shakey robot, 10–11
- Sigma, cognitive architecture, 200
- SIGVerse simulator, 125, 138
- SimAgent Toolkit, 470
- SIMNOS robot simulator, 468
- Sims creatures, 60
- Simulator, review of robot, 125, 131–137
- Situated cognition, 49, 217, 396, 398
- Situation calculus, 461
- Skill Transfer, manipulation, 326–327
- Skin, 145–164
- SLAM, 127, 137, 295–296, 305, 308. *See also*
 - RatSLAM
 - cognitive map building, 300–306
- SNARC effect, 442
- Soar, cognitive architecture, 182, 194, 199, 207, 444–445
- Social cognition, 14, 80, 200, 226, 277–278, 361–378, 391
 - psychology/neuroscience, 362–365
 - unified computational theory, 370–372
 - vision, 277–281
- Social learning instinct, 49–50
- Social robotics, 126, 379–383. *See also* Human-robot interaction
- Soft materials, 100–101
- Soft robotics, 12, 13, 99–120, 145, 151, 19
- Soft sensing, bioinspired, 109
- Sparse rewards, 259–260
- Speech. *See* Language
- Speech and vision, 281, 283–286
- SpikeStream simulator, 468, 470
- Spike timing-dependent plasticity (STDP), 153
- Spiking neural networks, 9, 30–34
- Spiking wavefront propagation, 31–36
- Stages, Piaget, 52–53. *See also* Development, nonlinear stages
- Stigmergy, 79–80, 93
- Superintelligence, 242
- Superorganism, swarm, 80, 91–93
- Swarm cognition, 87–94
- Swarm intelligence, 13, 77, 80, 94
- Swarm robotics, 8, 12–13, 65–66, 77–98
- Symbol emergence, 408
- Symbol grounding, 170, 220, 397, 442
- Symbolic knowledge representation, 418–423.
 - See also* Knowledge representation
- Syntactic, language analysis level, 397–398
- Synthetic methodologies, 9–10
- Tactile sensors, types, 168
- Teaching by demonstration (TbD), 317, 325, 332
- Teleoperated robots, 352–353
- Teleoperation, manipulation, 317, 323–324
- Temporal difference learning, 199
- Tests for Robot Consciousness, 467–469
- Time-delay neural network (TDNN), 176
- Tortoise, Walter, 4, 9–11, 19
- Total Turing test, 467
- Touch (tactile), 145–159
- Transducers, 146
- Trustworthy AI, 232, 351
- Turing test, 464–468
- Turtlebot robot, 125
- Uncanny valley, 381–382
- Universal gripper, 100, 109
- USARSim, 132
- U-shaped learning, 53
- Value systems, 25–26, 36, 50, 195
- Vehicles, Braitenberg, 7, 9–10, 19–20
- Virtual reality, 138
- Vision, 271–294. *See also* Perception
 - knowledge-based, 285
 - principles, 274–287
- Vision and speech, 281, 283–286
- VizDoom game, 69
- V-REP (CoppeliaSim) simulator, 125, 136
- Walking, 30, 48, 51, 60–61, 67, 99, 108, 129, 162, 279
- Webots simulator, 125, 134–136
- Wiskerbot robot, 22
- Word combination, 401
- Words as Social Tools, 447
- XCR-1 robot, 465
- YARP middleware, 126–127, 134

Intelligent Robotics and Autonomous Agents

Edited by Ronald C. Arkin

Billard, Aude, Sina Mirrazavi, and Nadia Figueroa, *Learning for Adaptive and Reactive Robot Control*

Dorigo, Marco, and Marco Colombetti, *Robot Shaping: An Experiment in Behavior Engineering*

Arkin, Ronald C., *Behavior-Based Robotics*

Stone, Peter, *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer*

Wooldridge, Michael, *Reasoning about Rational Agents*

Murphy, Robin R., *Introduction to AI Robotics*

Mason, Matthew T., *Mechanics of Robotic Manipulation*

Kraus, Sarit, *Strategic Negotiation in Multiagent Environments*

Nolfi, Stefano, and Dario Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*

Siegwart, Roland, and Illah R. Nourbakhsh, *Introduction to Autonomous Mobile Robots*

Breazeal, Cynthia L., *Designing Sociable Robots*

Bekey, George A., *Autonomous Robots: From Biological Inspiration to Implementation and Control*

Choset, Howie, Kevin M. Lynch, Seth Hutchinson, George Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*

Thrun, Sebastian, Wolfram Burgard, and Dieter Fox, *Probabilistic Robotics*

Mataric, Maja J., *The Robotics Primer*

Wellman, Michael P., Amy Greenwald, and Peter Stone, *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*

Floreano, Dario, and Claudio Mattiussi, *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*

Sterling, Leon S., and Kuldar Taveter, *The Art of Agent-Oriented Modeling*

Stoy, Kasper, David Brandt, and David J. Christensen, *An Introduction to Self-Reconfigurable Robots*

Lin, Patrick, Keith Abney, and George A. Bekey, editors, *Robot Ethics: The Ethical and Social Implications of Robotics*

Weiss, Gerhard, editor, *Multiagent Systems*, second edition

Vargas, Patricia A., Ezequiel A. Di Paolo, Inman Harvey, and Phil Husbands, editors, *The Horizons of Evolutionary Robotics*

Murphy, Robin R., *Disaster Robotics*

Cangelosi, Angelo, and Matthew Schlesinger, *Developmental Robotics: From Babies to Robots*

Everett, H. R., *Unmanned Systems of World Wars I and II*

Sitti, Metin, *Mobile Microrobotics*

Murphy, Robin R., *Introduction to AI Robotics*, second edition

Gruppen, Roderic A., *The Developmental Organization of Dexterous Robot Behavior*

Boissier, Olivier, Rafael H. Bordini, Jomi F. Hübner, and Alessandro Ricci, *Multi-Agent Oriented Programming*

Cangelosi, Angelo, and Minoru Asada, *Cognitive Robotics*