Embodying Design

An Applied Science of Radical Embodied Cognition

Christopher Baber

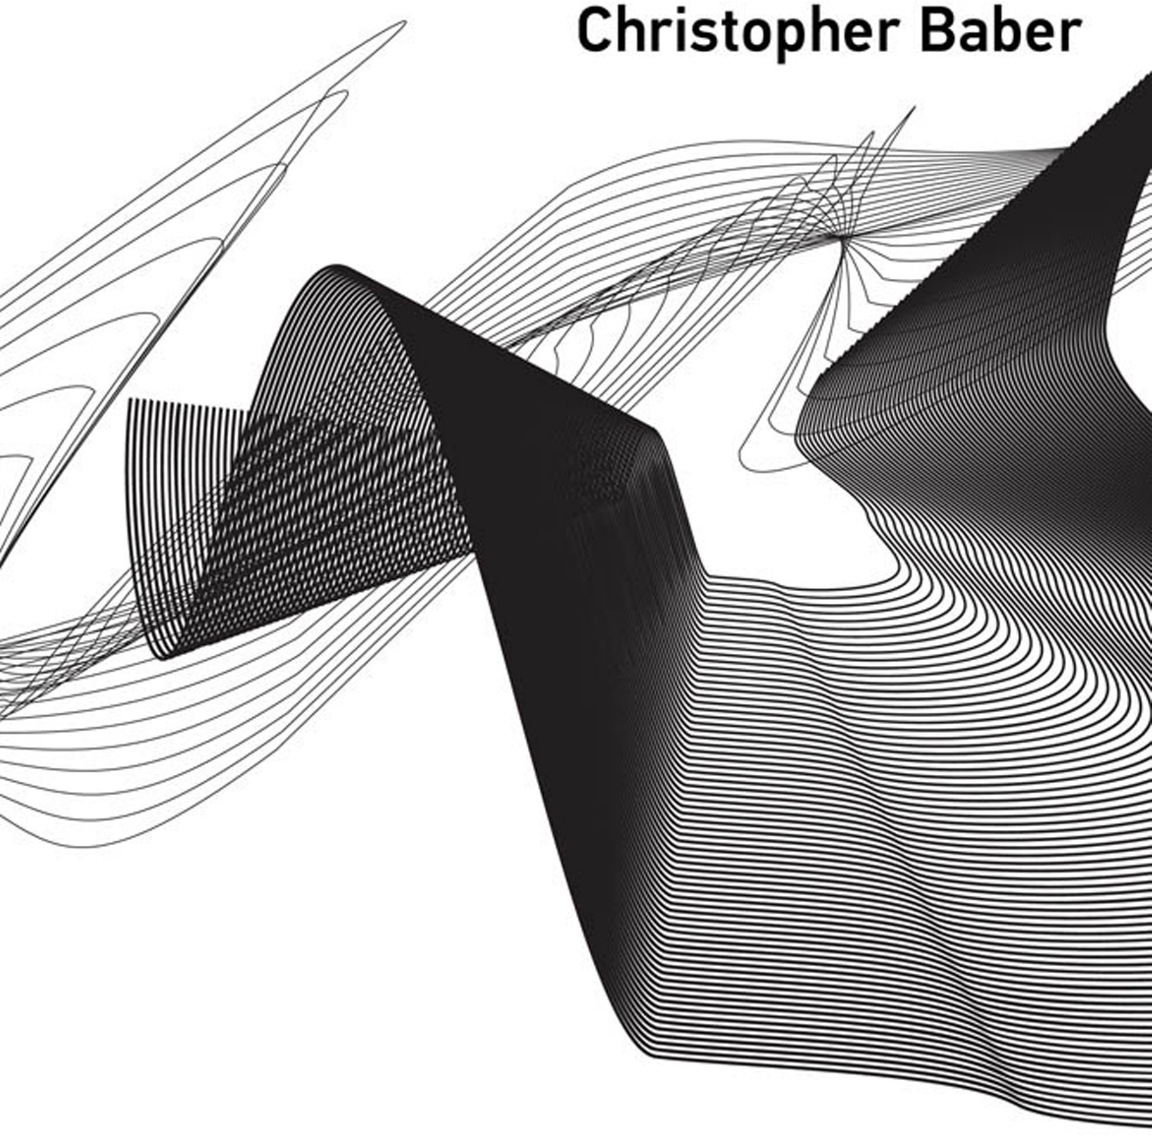

Embodying Design

Embodying Design

An Applied Science of Radical Embodied Cognition

Christopher Baber

The MIT Press
Cambridge, Massachusetts
London, England

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 Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Baber, Christopher, 1964– author.
Title: Embodying design : an applied science of radical embodied cognition / Christopher Baber.
Description: Cambridge, Massachusetts : The MIT Press, [2021] | Includes bibliographical references and index.
Identifiers: LCCN 2021033926 | ISBN 9780262543781 (paperback)
Subjects: LCSH: Expert systems (Computer science) | Human-machine systems. | Thought and thinking. | Artificial intelligence.
Classification: LCC QA76.76.E95 B22 2021 | DDC 006.3/3—dc23
LC record available at https://lccn.loc.gov/2021033926

In memory of
my brother, Jon,
&
for
my grandson, Arthur.

Our actions depend on finding their objects
And growing around them
Until one or the other is forced to bloom.

Douglas Crase, 1981, The House at Sagg, *The Revisionist,* Boston, MA: Little Brown and Co.

I'm painting, I'm painting again!
. . .
You can't see it 'til it's finished!
I don't have to prove . . . that I am creative!
. . .
All my pictures are confused!

Talking Heads, 1978, Artists Only, *More Songs About Buildings and Food,* New York: Sire Records

Contents

Preface

In writing this book, I am attempting to pull together many threads that have been gathered during thirty plus years investigating what it is that people do with digital technology. Over this time, my enthusiasm for digital technologies has waxed and waned either as a result of the opportunities I have had to tinker with devices or as the result of concerns over the reach that technology (and the organizations that control it) has into our everyday lives. In equal measure, this book reflects my journey from understanding human behavior in terms of cognitive psychology (as "information processing") to an appreciation of the significance of embodied cognition. Specifically, in this book, I employ my understanding of Anthony Chemero's radical embodied cognitive science (RECS) to some of the design challenges that digital technologies present.

My journey from information processing to RECS has involved a number of fortunate, often accidental, meetings, and I have benefited hugely from the opportunities that these have provided me. To put these meetings into some semblance of order would imply an organizing principle that is only vaguely correct, but this helps in telling the story of this book and how it developed. My academic career began in the applied psychology unit at Aston University, where I completed a PhD on the human factors of speech recognition, under the supervision of Rob Stammers and Dave Usher, in the late 1980s. This work explored the potential for speech technology to be used in the control rooms of electricity-generating power stations.[1] It was here that I learned about ergonomics and the delights of studying people doing their real work in their real work environments. I also learned about the perils and pitfalls of getting digital technology to behave in ways that would be beneficial, particularly the early forms of speech technology at our disposal. It was at Aston that I struck up a lifelong friendship and working collaboration with Neville Stanton, and some of the ideas that have germinated from our early work (particularly the notion of "rewritable routines"[2]) have a ghostly presence here. Over the intervening years, making sense of these ideas (and what it means to study people at work) has led me away from the "standard" information-processing approaches that informs so much of ergonomics to search for alternative theories and explanations.

During my PhD, I became interested in how people could use speech technology when they were outside the control rooms of electricity-generating power stations. Initially this involved a laptop in a rucksack with a small head-mounted Phase Alternating Line (PAL) television screen (so that the person had visual confirmation of what the computer had recognized). Over a few years, my research team and I developed wearable computers for maintenance workers, emergency services personnel, and crime scene examiners (CSEs). For the most part, the work was a mixture of hardware/software development with experiments and metrics to evaluate the impact of these technologies on people. From the work with CSEs, I began to think about sensemaking at crime scenes. This led to thinking about the ways in which concepts from distributed cognition could be applied to crime scene examination. In parallel with this, I was working with Neville on projects involving distributed situation awareness. The crime scene work led to two unexpected invitations. The first was to present the work to the Naturalistic Decision Making (NDM) conference, where I first met Gary Klein and Robert Hoffman. The idea that expertise can only be studied in "ecologically valid" settings, which this community strongly endorses, is central to my thinking. This is one of the reasons why I went to the UK College of Policing's Harperly Hall to study experienced CSEs and why I later worked with simulated crime scenes in Teesside to compare how experienced and trainee CSEs conducted searches.[3] The second was an invitation to attend the fledgling Distributed Thinking Symposium series that Fred Vallee-Tourangeau and Stephen Cowley ran from Kingston University. These symposia not only introduced me to the notions of interactivity but also to David Kirsh, Anthony Chemero, and Lambros Malafouris. Subsequently, the Distributed Thinking Symposium moved (with Stephen) to the University of Southern Denmark, where I met Christian Mosbæk Johannessen, who initiated an interdisciplinary project on writing and drawing, bringing together Marieke Longcamp, Susan Stuart, Paul Thiobault, and me.[4]

In an attempt to consolidate my ideas about the role of physical objects in CSE, I started to look at the ways in which people used tools. The literature seemed, back in the 1990s, quite sparse and spread across different disciplines that tended to have little connection with each other. So, I pulled together what I could find and wrote a book.[5] In part, this book was an attempt to make sense of tool-mediated interactions with the environment. I had benefited from discussions with colleagues at the University of Birmingham, particularly Ted Megaw (who had worked on ergonomics and motor control in the 1970s) and Alan Wing (who continues to define the field of how people coordinate physical movement). Both of them have an approach that marries engineering concepts (inspired by versions of control theory) with fundamental understanding of human activity, and both set up experiments that abstract the core features of real-life activity into tasks that are amenable to experimentation. While neither fully subscribed to the dynamic systems or RECS approaches in this book, I learned a great deal from them in terms of what a rigorous and testable description of activity ought to look like.

As I was writing *Cognition and Tool Use*, my thinking (while incorporating some aspects of distributed cognition and interactivity) was still influenced by information-processing concepts and the initial ideas of forms of engagement depended on “schema” and “automaticity.” I now recast the idea of forms of engagement to better fit with interactivity and embodiment, and the inspiration for this change has come from several sources. On the basis of the tool book, I was invited, by Witold Wachowski, to an AVANT[6] conference in Torun, Poland. Alan Costall, Robert K. Logan, David Kirsh, J. Kevin O’Regan, Richard Menary, Joanna Rączaszek-Leonardi, and Anthony Chemero were the other invited speakers. From this event, I was able to compare my own stumbling efforts to explain what people did with tools to more cleanly developed theories, particularly of David (in his account of how people use artifacts and actions to “do” cognition) and Tony (in his radical embodied cognitive science). The tool book also led to invitations from Lambros to workshops in Oxford to learn more about his material engagement theory, and from Blandine Bril in Paris to learn more about her theory of functional reasoning account of tool use. I have drawn heavily from all of these ideas and have attempted to find synergies and parallels between them, within the overarching framework that RECS offers. No doubt I am misinterpreting and twisting their arguments, but my misunderstandings are born purely from ignorance rather than malice, so I hope that they can forgive me. I urge the reader to go to the source material for these ideas. Lambros also encouraged Tom Wynn and Fred Coolidge to invite me to their workshops on applying material engagement theory to paleoarchaeology, where, alongside them and, among others, John Gower, Klint Janulis, and Leee Overmann, we discussed the nature of early hominid tool use.[7] At Birmingham, I have also benefited enormously from ongoing conversations with Andrew Howes on computational modeling of human decision-making.[8] More recently, Jan-Maarten Schraagen and Paul Ward, colleagues from the NDM conferences, invited me to contribute a paper on 4E (Embodied, Embedded, Enacted, Extended) cognition to their handbook on expertise.

I also want to thank Doug Sery and Noah Springer at the MIT Press for their help in taking this book from a sketchy manuscript to the version you are reading and to three anonymous reviewers, who have generously provided comprehensive and detailed reviews of the various versions of this book as it has evolved.

I am indebted to all of the people I have mentioned (and to the attendees of various workshops, symposia, and conferences and to all of the PhD students who have taught me through my supervision of them) for their inspiration and support in the development of the ideas in this book. In tracing the path from initial thinking (in distributed cognition and in making sense of how people use tools), it might appear as if there is a neat, linear path from "information processing" to "embodiment." I doubt that this is the case, and this book is, in part, a continued reorientation of my thinking from information processing to RECS as a way of explaining how people think and act. In particular, I have chosen to couple the consideration of digital technologies with a broader consideration of design and creativity partly because of ongoing discussions that I have had with Tony Chemero and partly because there seems to be a gap in the information-processing literature when it comes to creativity,[9] so it made sense to see how embodiment could plug that gap; and, of course, I liked the challenge of taking a theoretical position that many people dismiss as being about just "low-level" activity and demonstrating how it is equally applicable to high-level cognition, like creativity.

1 “Cut the Pie Any Way You Like, ‘Meanings’ Just Ain’t in the Head!”

Introduction

The title of this chapter quotes Hilary Putnam.[1] In a thought experiment, he asked the reader to consider twins living on different versions of Earth: in one, “water” had the properties with which we are familiar; in the other, “water” had different chemical properties but these properties could be described using the same words as used on our Earth. So, when twin 1 and twin 2 say, “Water is wet,” do they mean the same thing? For Putnam, the answer is “no” because, even though they are using the same words, the “truth conditions” (defined by the properties of the environments in which they live) create different contexts in which to interpret the words. Putnam’s quote can be repurposed as “cognition ain’t all in the head,” and this is a basic point that will be argued in this chapter.

I use embodied cognition as the lens through which to understand how designers engage in creative practices and also to understand how people use designed artifacts (in particular, digital technologies). In this respect, embodied cognition is playing a role in explicating design thinking (because “creativity” arises from interactions with materials rather than occurring solely in the head) and a role in informing design practice (by providing a theory of what people do with artifacts). Throughout the book, the phrase “embodied cognition” refers to the collection of theories that could be called “enactive,” “embedded,” “situated,” or “distributed.” I appreciate that my choice is controversial, but Shipp and Vallee-Tourangeau[2] point out that more papers use the term “embodied cognition” than the other terms. Depending on which review you chance upon, there may be three,[3] six,[4] or more flavors of “embodiment.” However, there is a broad consensus that humans, as cognitive agents, are *embedded* in environments in which they *enact* their *embodied* skillful coping in response to the scaffolding of artifacts that allow for the *distribution* or *extension* of cognitive activity.

Cognition and Embodiment

I spent many years working in the traditions of cognitive psychology, applying concepts and theories from this discipline to understanding people at work and their interactions with artifacts. Increasingly, I find that these concepts and theories are incomplete and do not capture the experience of either designing or using things. I believe that radical embodied cognitive science (RECS) provides a richer and more coherent account of what I find when observing and speaking to people in their workplaces or when evaluating prototypes than theories derived from cognitive psychology. Later in this chapter, I discuss RECS in more detail. For now, a quotation from William James, whose Principles of Psychology influenced not only cognitive psychology but also philosophy, particularly Pragmatism, illustrates the general tone of the argument.

> The world experienced comes at all times with our body at its center, center of vision, center or action, center of interest. Where the body is is “here”; when the body acts is “now”; what the body touches is “this”; all other things are “there” and “then” and “that.”[5]

Perhaps the word “embodiment” implies small children learning to count by using their fingers to represent the numbers 1 to 10. As an aside, the word “digital” is derived from the Latin for fingers (or toes). The use of the word “digits” to refer to numbers occurred around the fifteenth century, but it was not until the twentieth century that “digits” related to *all* numbers, and only in the last fifty years or so that “digital” came to apply to binary coding. More recently still, “digital” has come to apply to the technologies that make use of binary coding, with phrases such as “digital native” implying a facility with computer technology. So, in everyday parlance “digital” relates to fingers, to numbers, to technologies, and to the ways in which our information is codified. Information can be captured, processed, stored, and transmitted in digital form, and this is not simply a consequence of technology but is at the root of the “information-processing” models of cognition. It is against the broad concept of cognition as information-processing that theories of embodiment rail. Metaphorically, we might look for ways in which the original meaning of digital relates to both information and cognition.

Returning to the child counting on fingers; at one level, children associate their fingers with numbers. Very small children can recite the numbers 1 to 5 while touching their fingers, but this is not the same as knowing how to count[6] (ask a basic question about adding or subtracting, and they might struggle to answer). What the child needs to learn is the purpose of counting. Beyond a certain age, children might dispense with counting with their fingers and develop the ability to perform calculations "mentally." For Vygotsky, the crucial turning point comes from internalizing "rules" that apply to counting. This raises the question of what is being "internalized" as these "rules" are learned. For some writers, "internalization" merely means substituting the fingers on the hand for symbols in the brain.[7] One of the central debates (between "mainstream" cognitive science and embodied cognition) concerns this question of "internalization." In the version of embodied cognition followed in this book, human cognition can be explained without recourse to "internal representation."[8] It is important to note that this claim is not simply a matter of faith but requires a particular stance to research (both theoretical and methodological) that would allow us to define and demonstrate ways of explaining behavior that do not rest on internal representations. This position not only challenges basic assumptions of cognitive science but also, I argue, provides a richer and more parsimonious account of how people interact with artifacts and what designers do when they design these artifacts.

What Is Cognition, If It Is Not Information Processing?

The simple dichotomy between physical and cognitive activity implied by "internalization" misses essential aspects of the development of mathematical skills. Take the problem of solving simultaneous equations—that is, finding values for x and y that satisfy pairs of equations such $3x+y=11$ and $2x+y=8$. Several strategies can be applied to such problems. One approach, using elimination, recognizes that both equations have the same value for y (and if they do not, then it might be possible to manipulate either x or y, through multiplication or division, to make the values the same in each equation). From this, the solution involves subtracting one equation from the other (to find that, in this case, $x=3$ and $y=2$). Or you could plot a graph of these equations and find where the lines intercept. For both approaches, once you have learned the routine, solving the equations is a matter of applying the steps in a routine rather than "internalizing" any of the information. You might accept this point but argue that the steps are internalized. However, often the steps reframe the problem. That is, the experienced mathematician would either "see" the solution or "automatically" work through the steps until a solution was found. In this example, the information is the mathematical symbols, and the processing consists of the steps through which these symbols are transformed (together with an appreciation of when to stop transforming, i.e., what defines a solution to this problem). What the experienced mathematician develops is a way of defining the key features that are relevant to a problem and a set of actions that corresponds to these features.

From the example of solving simultaneous equations, we might ask what does cognition involve? In these examples, I have argued that "cognition" could be performed not in the head but through the manipulation of "external" information. So, what definition of cognition could allow both types of activity? At a minimum, cognition involves processes that can enable interpretation of salient information, coordinate actions on this information, judge the outcome of these actions and anticipate whether a given action is likely to be effective, adapt actions to increase the likelihood of effectiveness, and learn (or retain) effective actions.

To appreciate the depth of embodied cognition as a critique of information processing, we should immediately dismiss the suggestion that "embodiment" merely means "having a body." Some of the work relating to embodiment involves studies that make literal use of the word "body" and suggest that changes of the body, such as altering posture, can have a bearing on behavior. I am not convinced by such research as if often fails replication tests, so will not include it here. Alternatively, embodiment might suggest that there are some physical actions that we do during cognition, such as counting on our fingers. From the information-processing perspective, such actions are dismissed as incidental and as having no impact on cognition; the assumption seems to be that anything outside the brain (or anything that is not encapsulated in symbols) must relate to something other than cognition. The defining features of cognition I presented earlier do not demand either symbols or information processing. For embodied cognition, action *is* cognition.

A further problem with the claim that "embodied" means merely having a body is that it replaces the mind-body dualism of information processing with a body-environment dualism. For the theory of embodied cognition pursued in this book, it is important to recognize that the environment is integrated into cognitive processes. The boundaries between the components of the human-artifact-environment system (figure 1.1) are permeable. Obviously, this does not mean that artifacts will seep into the skin. But nor, I think, does it mean that the artifact becomes a part of the person.

Given the close coupling in the human-artifact-environment system, it becomes difficult (if not impossible) to claim that the elements of this system can be treated in isolation. This raises a question of where there are borders and boundaries in the system. For Sennett,[9] a boundary is an edge where one thing ends and another begins, while a border is a site of exchange. Recognizing the importance of boundaries, we can note that an artifact, such as a tool, does not become a "part" or an "extension" of the person (much as this has been proposed in discussion of tool use). Recognizing the importance of borders, we can appreciate how the artifact's functions will be modified by the person and the person's capabilities will be mediated by the use of the artifact; this is not due to the person becoming cognitively or physically enriched but rather due to the system having a new equilibrium. In other words, "in no system which shows mental characteristics can any part have unilateral control over the whole"; that is, "the mental characteristics of the system are immanent, not in some part, but in the system as a whole."[10] From this, the artifact offers new borders (between

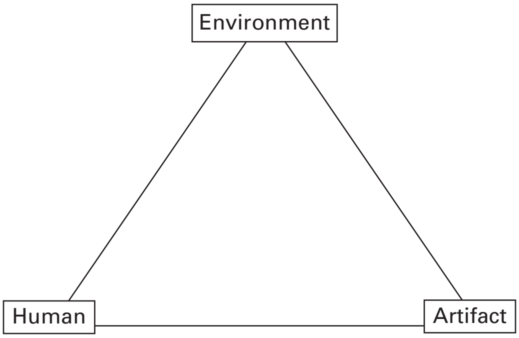

Figure 1.1
Interacting elements of a human-artifact-environment system.

person and artifact, artifact and environment) that create opportunities for the exchange of information and action. That is, cognition arises from the interactions between body, artifact, and environment.

We can't go much further in our discussion without addressing the question of how and why embodied cognition challenges the notion of representation. Indeed, a defining feature of much of the debate surrounding embodied cognition is the depth of anger and irritation that surrounds the very idea that we can dispense with the notion of representation. This debate has been characterized as the "representation wars."[11] Before wading into the debate, I note Dietrich's wry observation that "no scientist knows how representations represent."[12]

What Is Wrong with "Mental Models"?

"Internal representation" is the defining feature of information processing. For the information-processing view, the organism uses its senses to sample the environment. The resulting data are then translated into symbols that define meaning. This requires an appropriate apparatus to translate information from the senses into symbols and to process these symbols to create meaning. From this perspective, an "internal representation" is simply the side effect of using such apparatus—in other words, the symbols need to be put somewhere and they need a production line that manages their translation from sense data to meaning to physical action, with each stage of the production line performing a different operation on the symbols.

Interestingly, while the information-processing approach might *imply* the manipulation of symbols as a "language of thought," many theories developed within this tradition use different abstractions. For example, Baddeley's model of working memory[13] does not propose that we have a temporary storage of a symbols, such as words (e.g., when we remember a telephone number), but rather that data are stored in terms of temporal duration. That is, the "articulatory loop" (or phonological loop) has a duration of around two seconds and, like an old-fashioned tape loop, has new information overwrite existing information. Other notions of working memory (particularly the discussion of this concept in textbooks on human-computer interaction) assume that memory has a capacity defined by the quantity of symbols it can hold (e.g., 7 ± 2, derived from an experiment reported in 1956[14]). The reason for mentioning this is that the latter assumes that capacity (of working memory) is defined in terms of symbols, while the former assumes that capacity is defined by enaction—in this case, the time it takes to speak words. Indeed, there is a good evidence that working memory capacity for longer words differs from that for shorter words (which a "symbolic" account would struggle to explain but is obvious from an "articulatory loop" perspective).[15] These temporal dimensions of memory suggest that not all cognition involves the specification, translation, or manipulation of symbols.

The focus of the information-processing approach on thought as the algorithmic manipulation of symbols separates the thinking mind from the world that it occupies. The argument for embodied cognition is that, taking this point to its logical conclusion, *none* of what we have defined as cognition requires the use of such symbols (in much the same way that the examples of solving simultaneous equations by manipulating the printed symbols does not require these symbols to be internalized). This would mean that, to use Chemero's phrase, the "mental gymnastics" required in information processing (e.g., in terms of translating between environmental information and mental symbols) is not necessary.

I find the term "internal representation" confusing, so I am going to use "mental model" instead (on the assumption that this describes a "model" of the environment that is created and stored in the mind). An information-processing view of human cognition assumes the representation of information, extracted from the environment, in the form of symbols. These symbols are defined by structural units that are either "word-sized concepts"[16] or "icons,"[17] and cognition involves the manipulation (according to specified rules) of these symbols.[18] As noted previously, information-processing approaches assume some apparatus that performs translations of features of the environment onto internal states, which can result in the ability to act on the environment. For embodied cognition, we might ask what this "apparatus" might be (if not information-processing apparatus in the brain), what this "information" might be (if not symbols), and what the internal state might be (if not a mental model)?

When the phrase "mental model" is used, the same collection of authors tends to be cited in its support, one of whom is Kenneth Craik. Here is a quotation of his that is commonly used:

> If the organism carries a "small scale model" of external reality . . . within its head, it is able to . . . react to future situations before they arise, utilize the knowledge of past events in dealing with the present and future, and in every way react in a much fuller, safer, and more competent manner to the emergencies that face it.[19]

Craik's idea satisfies the requirements for cognition outlined earlier (in terms of anticipation, learning, and effective response). A loose reading of this quotation would find little to separate the idea of a "'small scale model' of external reality" from a "word-sized concept" or "icon" or "mental model." But this is *not* what Craik is arguing for. Prior to this quotation, Craik used the Kelvin tide estimator as an analogy for his concept (and had also written papers using servomechanisms to model [human] radar operators[20]). But, this is markedly different from the idea of a mental model that, say, Frederic Bartlett[21] assumed in his discussion of the gist of stories. For Bartlett, a mental model is a summary of salient information from which to build interpretations, judgments, decisions, and actions. For Craik, the "small scale model" had to be a "physical working model" that "shares a relation-structure to that of the process it imitates."[22] What is important here is that he is *not* claiming a mental model that represents reality but a *process* that mirrors reality. Craik's thinking was, to some extent, influenced by the UK cybernetics movement in the 1940s, and in particular by the work of William Grey-Walter, pioneer of robots as autonomous entities. In one visit, Craik and Grey-Walter discussed "the aiming accuracy of air gunners" and how the activity could be explained in terms of "goal-seeking and scanning. . . ."[23] The resulting mechanical conception reflects embodied cognition's notion of perception-action coupling.[24] That is, how our "lived body"[25] "opens the world to us as full of possibilities for action."[26] Features of the environment are perceived, and these features are associated with action. Key to this proposal is that there is no requirement for the features to be translated into symbols. Rather, the perception of features is direct. For me, perception-action coupling defines the relation structure that Craik is discussing. As we grow from baby to toddler, the range of possibilities for action increases. The relationship between action and the environment can be considered in terms of Ashby's law of requisite variety. This law states (in cybernetics terms) that a "controller" can model the environment that it is controlling only if it has enough variety to respond to the states that the environment exhibits; if the environment becomes more complex, then the "controller" needs to create new models or else its uncertainty increases.

The question is whether this “model” is a replica of the environment built in the mind. Or whether it is a repertoire of action to allow you to act effectively in and on the environment. The environment can be responded to through the mechanism of a “physical working model.” Applied to cognition, we could say, to use Rorty’s lovely phrase, the brain is for “coping, not copying.”[27] In other words, the brain is for acting rather than representing. People develop strategies that allow them to use features in the environment to structure cognitive activity.[28] The environment *is* the representation of a problem; in other words, the “external” information in the example of solving simultaneous equation is the “environment” for this particular activity.

The very notion that the external environment needs to be represented as a mental model in order for the person to perform an action requires the assumption that the person and the environment are not only physically distinct but also cognitively separate. For embodied cognition, the person and environment are mutual (i.e., linked in a way that one implies the other) and reciprocal (i.e., linked in a way that one affects the other) and form a self-organizing system. Thus, the behavior of a human-artifact-environment system involves continual adaptation as it self-organizes. From this, cognition is “a kind of dynamic adjustment process in which the brain as part of and along with the larger organism, settles into the right kind of attunement with the environment—an environment that is physical but also social and cultural.”[29]

From an information-processing perspective, the construction and use of mental models come with processing costs. There is clearly a significant degree of mental effort involved in constructing or learning a mental model. The pay-off is assumed to be that once this is built, it can be reused and hence the effort is an investment. But this assumes that the mental model will be generalizable. There are several problems with this assumption. The first is that what is learned in one situation might not be appropriate to other situations. If the situations *were* constant, there might be much simpler means of capturing their essential aspect to ensure a consistent response. For embodied cognition, this “essential aspect” involves perception-action coupling, or physical action in response to features in a given situation. Repeated exposure to this situation increases the probability of the action. Of course, a “mental model” might reflect the essential aspects in just enough detail to provide flexibility for future situations.

But, if this were the case, there would be little need for a "model" as such. Rather, recognition of these essential features ought to be sufficient. This suggests that information processing requires *another* set of symbols (to be stored in their own form of memory) that correspond to these essential features. For embodied cognition, the features simply exist in the environment and, in a very real sense, are "stored" there.

A second problem with mental models is that there needs to be more "rules" that encapsulate the knowledge of how to respond to these features. This "know how"[30] or "tacit knowledge"[31] relates to our skillful coping with our environment. Information-processing approaches tend to use the clumsy argument that knowledge is either "declarative" (i.e., facts, or propositions, represented in some form of symbolic language) or "procedural" (i.e., anything to do with activity)—which allows "procedural" knowledge to either be magicked away or be subsumed under the aegis of declarative, such that symbolic information takes the form of "production rules" (i.e., if condition *x*, then action *y*). A third problem is that symbols require a "semantics," such that they can be labeled in terms of their salience. But, to assign a meaning to symbols requires a further set of symbols (with the requisite information-processing apparatus). A fourth problem is that the sort of content specified by the symbols used by information processing ought to allow us to make judgments over its quality. This has been called the "hard problem of content" and is well expressed in the following:

> Anything that deserves to be called content has special properties—e.g., truth, reference, implication—that make it logically distinct from, and not reducible to, mere covariance relations holding between states of affairs.[32]

For a mental model (or any other form of internal representation) to have scientific credibility, it needs to be something that has a substantive role in cognition. I am not sure that even people who study mental models believe that these are anything other than convenient fictions. There is general agreement that mental models are incomplete, imprecise, ambiguous, fuzzy, poorly organized.[33] Even if mental models existed, there would need to be some further "perceptual" process by which these were interpreted—which implies the oft-parodied inner homunculus.[34] If we dispense with a homunculus to observe the mental model, there remains the question of how the mental model can have an impact on our actions. One conventional argument is that, having constructed a mental model, the brain then constructs a program (or set of instructions) that is passed to the body in order for it to act on the world. But this assumes an information processing apparatus that, oddly, does not include the body. For Gallagher, the information-processing approach is promulgated by "body snatchers."[35]

Why would the human brain put effort into constructing a representation of the environment, then analyzing this representation, then planning an action based on this representation, then simulating the outcome of this action by running it through the representation . . . all before acting? Put simply, if there is a thing in front of you, why would it be necessary to create a representation of this thing in order to pick it up? While this argument might not immediately explain how we can think about things when they are not in front of us, it allows us to wonder what alternative to an information-processing account could be offered. In order to consider this, it is necessary to reconsider what we mean by "information."

What Information Is Being Processed?

How do we make sense of an artifact? An obvious answer is to say that we obtain information from it. But this does not tell us what we might mean by "information." In one sense, information is a digital code (in the form of binary digits, or bits) that allows a computing device to run operations on data (the digital code is used to describe both the operations and the data), and this digital code defines the on and off states of transistors. In its earliest inception, the information-processing approach used the computer as a metaphor for the brain: both had input (in the form of data) that was manipulated (in the form of symbols) to produce output. For some early writers in the information-processing tradition, neurons in the brain behaved like transistors, switching on and off as information passes through them; but this rested on a whole bunch of assumptions which are manifestly untrue of the electrochemical activity of the brain. In the cybernetic tradition preceded information-processing view of cognition, switching related to control mechanisms that aligned action to environment. Some of these ideas reappear in various guises in theories explored in this book. Given that the metaphor does not apply to the workings of the apparatus, does it apply to the "stuff" that is being processed? Digital information is clearly not "information" for you or me when we are picking up a very full cup of hot coffee. So, what is the information we obtain from a cup?

The language used to describe our interaction with artifacts is problematic. The division of these interactions into subject (person) and object (artifact) means that these can be seen as discrete entities. In this way, there is a linguistic division between subject and object. Given this apparent separation, it then becomes necessary to introduce additional processes that can bridge this. Hence, information-processing approaches to cognition introduce discrete stages in which information is translated, such as input, processing, and output, as well as a separation of actions into discrete stages, such as begin movement, reach to object, pick up object. In embodied cognition approaches, making these distinctions is pointless because the system would always be in flux as it self-organizes in response to the disturbances caused by each element. Indeed, for embodied cognition, the division between subject and object becomes irrelevant; there can only be a "system" in which human and artifact join together (in an environment). As Samuel Butler has it, "Strictly speaking, nothing is a tool except during use."[36] From this, the joining together, implied by the term "use," creates a balance of activity between human and artifact (with both responding to their environment) in which they are mutually responding to the actions and effects of each other. For Varela's enactivist account, "in-formation appears nowhere except in relative interlock between the describer, the unity, and its interactions."[37] From this I infer that the human-artifact-environment system creates the unity within which, through its interactions, information is created.

I find it useful to distinguish between information-as-content (which requires processing) and information-as-context (which constrains action). One reason why information processing relies on a mental model of the environment is that it is supposed to allow the person to make predictions prior to performing an action, which reduces reliance on feedback from the environment. The argument is that such feedback can be time consuming, particularly if the person is processing this in incremental stages during the performance of an action. A further justification the information-processing approach offers for mental models is the "poverty of the stimulus."[38] This assumes that the environment rarely contains fully specified details for information processing, so the information-processing apparatus needs to supplement sense data. For me, this argument puts the cart before the horse; only if you assume that this apparatus is used to build a mental model is sense data insufficient. If we return to our over-full coffee cup, do we need to define the content of this scenario in order to guide our action? By describing it verbally, I have, of course, provided information-as-content (cup, liquid, temperature, capacity, spillage, scalding, and so on), and this might be one reason why it is so easy to assume that words (and other symbols) must also be the language that the brain uses to engage in cognition.

Around the same time that digital computers were developing (by which I mean sometime in the 1940s), Shannon[39] was developing information theory. For Shannon, the purpose of information was to reduce ambiguity in a message. Rather than consider the "meaning" of a message, he described it in terms of ambiguity: as the number of message elements increases, so the message can be more ambiguous (or, in his terms, have higher entropy, or disorder). So, the purpose of information in this view is to help maintain order in the transmission of messages. If you consider a math problem from high school, say, the probability of drawing a blue marble from a bag of mixed colors, the number of marbles you need to draw out (or the number of "questions" you need to ask) is determined by the context (i.e., the number of alternatives) and *not* the content (i.e., the example works whatever combination of colors or objects or containers we use). For information theory, then, the purpose of information is to reduce uncertainty by providing context. The units of information in this case can be thought of as "yes" or "no" and will be represented as binary digits, or bits (as an aside, for information theory, the bits have no meaning other than their role in managing uncertainty, while in computing the bits have the unique definition of a program instruction or alphanumeric character).

In information theory, Shannon defines uncertainty, or entropy, in terms of the probability of features in a set; sets of features that have low entropy are predictable (due to their low variability), while sets of features with high entropy are much harder to predict. Information, from this perspective, can be defined only with reference to something else; it cannot be defined independently but only in terms of difference. From this, we can think of an environment in terms of degrees of freedom (defined by the features and their possible combinations). While information theory would have been familiar to Gibson, he did not apply it in his ideas of how features of an environment support action.[40] "The term *information* cannot have its familiar dictionary meaning of knowledge communicated to a receiver. This is unfortunate, and I would use another term if I could."[41] Often the use of the word "information" caused Gibson problems because he wanted it to mean, at different points in his argument, a vehicle of communication, a form of knowledge, culturally modified content, and a naturally occurring ("invariant") property of the environment. In embodied cognition, the organism uses its senses to collect information from the environment, defining those features against which actions are possible, that is, information provides context. This does not mean that humans are unable to respond to content, just that this is not necessarily part of everyday cognition.

For this book, I will use the phrase "ongoing, reciprocal engagement" to reflect the enactive nature of the routine skills that involve "skillful coping."[42] The idea that we require complex apparatus to process information extracted from the environment (as per the information-as-content approach) commits us to viewing the brain as sluggish, clumsy, and poorly adapted. For embodied cognition, actions are guided by salient cues from the environment, and meaning is defined in terms of the consequences of action. The organism performs an action and the state of the environment changes. If this new state is acceptable, action stops, or the organism repeats the cycle of sampling and acting.

The information-processing approach presupposes that the organism's intent is a well-defined representation of the desired state of the environment. The embodied cognition approach presupposes that the organism has no "model" to aim for (although it does imply some criterion for acceptability). The first view assumes that perception (i.e., processing information from the senses) has the aim of constructing a representation of the organism's environment. The second assumes that perception is for action. In other words, the views can be distinguished by their focus on "world-in-the-mind" versus "mind-in-the-world."[43]

What Is the "Mark of the Cognitive"?

The distinction between "world-in-the-mind" versus "mind-in-the-world" can also be found in the field of distributed cognition, which emphasizes that humans use artifacts to "off-load" activity that is essentially cognitive[44]. For example, we use all manner of artifacts to help remember information (e.g., shopping lists, electronic diaries, the phonebook in our cell phone, and so on). We also use artifacts to perform manipulations on information (e.g., abacus, slide-rule, calculator). In distributed cognition[45] artifacts are "external representations" that become part of an information-processing system. This is also related to enactivist approaches in which the environment helps to structure problem-solving.[46]

In his account of calculating speed on a US Navy ship, Hutchins[47] discusses how several people perform tasks that contribute to sighting landmarks, making timings, marking a chart, and so on. In his study of medical records, Nemeth[48] shows how the physical attributes of the files on the end of beds in hospitals can tell a lot about the patient—for example, the number of pages or how creased or folded they are can tell how long the patient has been in the hospital, how many tests have been administered, and so how complex the case might be. However, for much of the distributed cognition literature, "cognition" is being done in the head of a cognizer (i.e., the human) rather than in the artifact. In this book, the argument is that cognition occurs in the *interaction* between person and artifact; as the person acts on the artifact to change its state, so the artifact provides opportunities for action (in a task-artifact cycle, see chapter 5) and also produces changes in the person. This raises the question of the extent which an artifact can participate in cognition.

Clark and Chalmers[49] use the example of Otto's Notebook to illustrate this claim that our cognition extends into our objects. In this example, Otto has impaired memory and so relies on the notebook to store information that he might require, such as directions to buildings. This notebook is functionally equivalent to brain-based memory for Otto. In part, this is because losing the notebook would, for Otto, mean the loss of the knowledge it contained—as if Otto, in losing the notebook, had lost his memory. In other words, Otto's notebook is not simply a passive store of information but an active component in his cognitive system, so that loss or damage to it would be functionally equivalent to loss or damage to any other part of his cognitive system.

Adams and Aizawa[50] argued that a fundamental problem with Otto's notebook having a structural role in cognition is that it invokes a "coupling-constitution fallacy." For them having the notebook available to be consulted does not make this notebook part of any cognitive process. Rather, the "mark of the cognitive" can be defined as the nonderived content brought by a cognizer. By way of analogy, they draw on the well-worn example of the "white stick" that blind people use to aid their navigation; while the stick plays a role in navigation, the stick does not, itself, "know" anything about its environment any more than the notebook "knows" what the words it contains mean. What seems key to their idea of a "mark of the cognitive" is the capability to manage information-as-content, not simply in terms of obtaining information from the immediate environment (or artifacts within that environment) but also of combining it with other information known by the person. However, this position is at odds with the "loop between brain, body and technological environment"[51] that is inherent in the Otto notebook example. In order to access the knowledge held in the notebook, Otto needs both an awareness of what the notebook contains, a strategy for accessing this knowledge, and the motivation to perform such a strategy. To say that the information known by Otto is held in the notebook is no different from the off-loading of information that distributed cognition emphasizes (a contemporary analogy is the way that we use the "phonebook" in our cell phone to store contact details). But for Clark and Chalmers, Otto's notebook is not simply an artifact that allows the user to off-load information; it is the instantiation of what Otto knows. However, the argument rests on the belief that "information" is content and, as represented in the form of symbols, can be stored in the brain or in a notebook. In neither of the positions presented here do we see the embodied cognition idea that I have termed "information-as-context." To better appreciate this point, we should turn our attention to the different schools of thought that address embodied cognition.

Perspectives on Embodied Cognition

What if we had a theory that dispensed with the need to model the world and that removed the need for the apparatus of information processing? What if, as Brooks notes (from his work in robotics), "the world is its own best model. It is always exactly up to date. It always contains every detail there is to be known. The trick is to sense it appropriately and often enough."[52] If embodied cognition relies on physical engagement with the world around us (in order to "sense it appropriately and often enough"), we face several questions—not least of which is why would physical engagement be something that is *not* part of an information-processing account of cognition? Revisiting Otto's notebook as information-as-context we might say that the content becomes salient when Otto consults it, and that salience arises from the ways in which this consultation is performed. For example, Otto flicks through the notebook in search of content to support a specific query, such as where is the Museum of Modern Art.

A criticism commonly levelled at embodied cognition, whatever its type, is that it defines "action" in terms of physical movement (even it is quite complicated, as in catching a flying ball) and "decisions" in terms of choosing a small number of cues. In effect, the complaint is that embodied cognition has failed to engage with "representation hungry"[53] domains—that, instead, it engages with "domains in which suitable ambient environmental stimuli exist and can be pressed into service in place of internal representations."[54] In particular, the complaint focuses on the challenge of cognition that involves the "absent" (i.e., how, in the absence of the cues in the environment, does cognition operate?) or the "abstract" (i.e., how, in the absence of concrete cues, does cognition operate?). In other words, how can embodied cognition deal with complex cognitive behaviors such as design or creativity?

In the case of "absent" stimuli, embodied cognition *could* rely on the repetition of prior actions.[55] When we have performed an activity to effect an outcome that is satisfactory, the sequence of actions could be represented as perceptual symbols.[56] Here, perceptual symbols are neural traces arising from sensorimotor performance, and their activation can result in the performance of the sequence of actions. For me, this feels like symbolic representation, although clearly perceptual symbols are *not* a set of instructions so much as the trace memory of coordinated neuromuscular activation (not dissimilar in concept to the notion of mirror neurons[57]). As Dreyfus puts it, "Past experience has set up the neuron connections so that the current perceptual input, which is similar to some part but never exactly like it, puts the brain area that controls movement into a specific energy landscape."[58] While I can see the basis of this argument, my concern is that it is overly focused on a brain-bound perspective, which loses sight of interactions within the human-artifact-environment system. A complementary but different concept, "embodied intelligence,"[59] emphasizes the importance of "performative awareness"—which is the phenomenology of the movement of the body in action, particularly for the skillful practitioner.

We have well-organized ways of moving our bodies, as the result of our continued experience of moving around in a physical world. This means that not only do we form "chunks" of action in cognitive terms, but that firings of muscles occur together in physical terms. In his study of human movement, Bernstein[60] defined degrees of freedom (DoF) as the combination of movements that are possible with, say, each joint in the arm. In an action, such as reaching to pick up a cup of coffee, each of the joints in the human arm can move in a variety of ways (defined as their DoF), including flex, extend, rotate, and so on, and the combination of the DoF of each of the joints can result in many different ways to perform the same action. The fact that we tend to perform similar actions in similar ways suggests that the DoF problem has a solution that results in consistent movement.

For Gibson, "Locomotion and manipulation . . . are controlled not by the brain but by information. . . . Control lies in the animal-environment system . . . ; behavior is regular without being regulated.."[61] In this view, rather than assuming a "controller," "regulation" arises from the animal-environment system seeking stability and avoiding entropy. For Bernstein, repeated performance of the same movements reinforces the activation of specific muscles to move specific limbs, so that these form "coordinative structures," which are "macroscopic spatio-temporal patterns"[62] of musculoskeletal activations that simplify the DoF problem. While Bernstein focused on the musculoskeletal structures recruited in the performance of actions, a similar concept is proposed by Luria in his suggestion that repeated examples of a movement become imbued with "kinaesthetic melodies."[63]

Bernstein's notion of "dexterity" involves balancing between stability of these coordinative structures (in order to allow an action to be repeated) and adaptation (to cope with changes in environment or task demands). But while coordinative structures provide a neat explanation of how we are consistent in our movements, we also need to recognize how movement adapts to small changes in situational features. The challenge of explaining dexterity (as the balance between consistency and variability in movement control) relates to the proposal that embodiment is ongoing, reciprocal engagement (with its emphasis on adaptive coping with the changing environment). This highlights the tension between ensuring consistency of response while adapting to variability in the environment. For me, this trade-off (between consistency and variability) has to be considered in terms of the balancing of activity within the human-artifact-environment system. Sampling the features requires effort, so optimal performance would involve minimizing the entropy of the environment by continually minimizing its DoFs. This points to the need to discover ways of reducing variability (both in terms of sampling features, i.e., exploring, and acting on the environment, i.e., exploiting opportunities to act). However, it makes little sense to treat each situation as if it was novel. Rather, we need to find consistent ways to respond to similar situations. From this, an information-processing approach would argue that a mental model provides us with the ability to define and store those features that define "similar" situations. That is, information-processing approaches assume that consistent movement arises from a "controller" in the brain that sends commands to the joints in the form of a "program" (not unlike the software that a computer uses) that defines when, and to what extent, each joint moves.[64] In such approaches, information-as-content is used to specify the movements of each joint, the location of the object to pick up, the path that the hand will follow to grasp the object, and the properties of the object itself (mass, center of gravity, and so on). This "content" is constructed from sensory data to create a "mental model" from which the specific the motor program guides movement.

A fundamental aspect of embodied cognition approaches is the close coupling within the human-artifact-environment system: the person's actions change in response to the state of the artifact or the environment (and, of course, the person's actions change the state of the artifact, and the artifact will change the state of the environment). Even in this simple three-element system, the manner in which "change" occurs will vary. Some of these changes will lead to stability in the system. In such circumstances, the system is well ordered and said to be self-organizing (and this might be a desirable state; equally, in terms of errors and accidents, the state could be undesirable). In others, the changes lead to instability and the system becomes disordered.

The body is considered to be part of a larger cognitive system.[65] From this, the ways in which the body moves (e.g., gestures, changes in posture, mobility, and so on) have an influence on cognition. In broad terms, "the brain is not the sole cognitive resource we have available to us to solve problems. Our bodies and their perceptually guided motions through the world do much of the work required to achieve our goals, replacing the need for complex internal mental representations."[66]

Even when we are not physically engaged with the environment, cognition draws on sensorimotor activity.[67] As Lakoff and Johnson[68] point out, there are many common metaphors that draw on our understanding of how the world relates to the movement of our bodies and the actions that we perform. Metaphorically, ideas *are* objects, and the mind *is* a container for these objects; we speak of grasping a concept. In this respect, these “metaphors we live by” hint at some underlying appreciation that cognition and physical activity intertwine. For this school of thought, metaphors are not simply words and phrases we use, but indices of cognitive structures we have acquired through our physical interactions with the world.

Radical Embodied Cognitive Science

There are many varieties of embodied cognition, but my preference is for Chemero’s radical embodied cognitive science (RECS). Before we go further, a definition of the theory would be relevant, and I am taking this from Anthony Chemero:

> I hereby define radical embodied cognitive science as the scientific study of perception, cognition, and action as a necessarily embodied phenomenon, using explanatory tools that do not posit mental representations. It is cognitive science without mental gymnastics.[69]

RECS challenges the assumptions that cognition must involve symbolic representation and a mental model of the environment in order to produce action. This does not necessarily mean that there is no “representation.” As we noted, for features of the environment to be responded to, there is a need to have some form of “information,” which, in turn, requires some form of “interpretation.” The distinction is a not a matter of all or nothing so much as a contrasting of “action-oriented” and “objectivist” representations.[70] From an information-processing perspective, the question is whether “action-oriented” representations (which explain skillful coping through the use of coordinative structure and “kinaesthetic melodies”) can be considered to be “genuine” representations—but this seems to assume that a “representation” can take only the form of a mental model (or symbols that can be processed by information-processing apparatus) rather than that of a “mediating state.”

RECS combines the notion of perception-action coupling (specifically through Gibson’s notion of affordance which is discussed further in chapter 4) with methods and metrics from dynamic systems to explain how behavior occurs in the context of ongoing sequences of action, adapting to system constraints. Such metrics allow quantification of the behavior of loosely coupled systems and provide insight into the ways in which the behavior of such systems has to be considered in terms that do not allow individual elements to be separated from each other (which is one of the reasons that I feel the concept of "affordance" is so often misconstrued). Accepting that these systems are non-decomposable leads to two assertions. The first is that the environment is constitutive of the system and one cannot suggest a separation of environment from organism any more than one can suggest a separation of organism from task. The second is that, following the first, one can discount the "coupling-constitution fallacy,"[71] which implies that the organism, by virtue of being distinct from its environment, must create a representation of that environment in order to act upon it.

RECS provides an account of "cognition" not simply as the consequence of a "brain-in-action," but also in terms of solving problems, making decisions, and performing other actions that are characteristics of cognition. In other words, cognition relates to the coordinated and adaptive response of the organism to its environment in the pursuit of tasks and goals. An obvious issue arising from this final point concerns the source of "goals." If, as the preceding points might imply, the organism's activity occurs in the context of an environment that changes in response to previous actions, one could ask what initiates an action and (equally) when does an action achieve an acceptable outcome? Taken to its extreme, this question concerns whether RECS is able to account for those activities that do not have an obviously "embodied" element, such as invention or creation or imagining or dreaming. RECS has tended to focus on relatively prosaic activity, such as categorical perception or locomotion, primarily because the modeling required to describe these activities in terms of nonlinear dynamics is challenging. This means that much of foundational research on RECS has concentrated on activities that are, in a sense, only partially or minimally cognitive.

As I will explain in chapter 2, design and creativity need to be considered in dynamic rather than discrete terms. The initial mark an artist makes on the canvas or the initial centering of a wedge of clay on the potter's wheel constrain subsequent actions. The artist creates, and responds to, changes in the affording situation. But such an idea can extend to most activities that we call "cognitive."

2 Thinking, Acting, Creating

Introduction

While creativity might strike the reader as an archetypal form of "abstract" cognitive activity to be performed "in the head," it has been poorly served by the information-processing literature. Indeed, for many writers the study of creativity is currently in some sort of crisis.[1] In part, this might be because the concept of "creativity" is not always clearly defined and so cannot be reduced to something that can be amenable to laboratory experiments. Research on how we might understand "creativity" has reached an impasse, stuck in divergent thinking tasks,[2] such as "multiple uses of a brick."

In chapter 1, I proposed that the word "information" has (at least) two meanings: one related to the information's content and one related to the context in which the information is presented. In terms of design, theories that inform "creativity" (and, to a lesser extent, "design thinking") tend to focus on information-as-content and ignore information-as-context. Let's say that "content" relates to the form of the artifact and that "context" refers to the environment in which the artifact is being used. Many theories of design thinking focus on content and draw heavily from information-processing theories of cognition. Many design textbooks propose that design is a form of problem-solving. Indeed, the very suggestion that design is about problem-solving contains within it the implication that the artifact represents a "solution" to a specific problem and that this can be defined in terms of content (which can be conceptualized, manipulated, and communicated). For example, if the problem is how to contain hot liquid for drinking, one solution might be a tea cup. The notion of information-as-context implies the need to focus on the environment in which artifacts are used—that is, the ways in which the ambiguity of an artifact's "meaning" needs to be resolved in order to let people use it. This is a perspective that is inherent in design *practice* and can be seen in some versions of design thinking (although much of this discourse still draws heavily on information processing and problem-solving). Furthermore, focusing on information-as-context is what designers do as a matter of course, but theories based on information processing ignore or trivialize this.

In this chapter, and throughout this book, my arguments overlap with those of Kees Overbeeke,[3] who led the Designing Quality Interaction research group in Eindhoven. He long argued that designers give too much weight to "cognitive" (i.e., information-processing) skills at the expense of technical or craft skills. Drawing on the work of Gibson and Merleau-Ponty (as radical embodied cognitive science does), he emphasized how meaning emerges in interaction. In this chapter, my argument is simply that embodied cognition helps theorize context, as it is experienced by designers, as the ongoing, reciprocal engagement in a human-artifact-environment system engaged in high-level cognition—such as creativity.

Convergent and Divergent Thinking

The manner in which designers respond to what Charles Eames called "a willing embrace of constraints" has been explored by Peter Rowe in his influential (and very readable) book *Design Thinking*. Rowe is concerned with "the situational logic and the decision-making processes of designers' action, as well as with theoretical dimensions that both account for and inform this kind of undertaking."[4] Taking his cue from Simon's conception of design as problem-solving, Rowe proposes that design involves the ability to respond to problem-oriented constraints in ways that adapt to the "covering characteristics" (or specific circumstances) in which the designer is working. Using concepts drawn from a blend of information-processing theory and phenomenology (particularly ideas from Merleau-Ponty) he proposes that

> the design process may be seen to be marked by a sequence of episodes or situations that are, in turn, coincident with periods of heuristic reasoning through which problems are defined and solutions sought. During each episode a particular heuristic device or set of devices can be said to be in operation and in general control of the reorganization of a problem space. Further, the orientation of this

> operation is neither entirely objective nor entirely subjective. It is both. Between episodes, control is relinquished, so to speak, from one set of organizing principles to another.[5]

A "problem space" is a set of plausible solutions to a problem (given certain constraints, such as the "rules" by which a problem could be solved, the features of the problem available to the problem solver, the end-point or "goal" of the activity of problem-solving). While Rowe's arguments employ the language of information processing (e.g., "heuristic reasoning," "problem space") and speak to the idea of information-as-content presented in chapter 1, their tenor is more suited to the experience of design and the need to work with "organizing principles." For me, this feels less like an argument based on information-as-content (in which the designer perhaps builds a mental model from which to imagine design concepts) and more on the practical, physical interaction in situations in order to explore and respond to constraints (organizing principles). Rowe, following Merleau-Ponty, considers a "situation" as involving the focused attention of the problem solver (this is similar to Csikszentmihalyi's notion of "flow"[6]). Situations, from this perspective, are ambiguous not only in their open-endedness but also in the dependence on the prior experience of the person experiencing this situation, who may or may not have a sense of how to respond. The "situation" can be regarded as a "wicked" problem (which does not have an obvious solution). To complicate matters (for an information-processing approach), heuristic reasoning becomes less about the simple application of "rules of thumb" (as might be implied by a literal reading of information-processing concepts), and more a matter of relying on general principles that are "sedimented" (using Merleau-Ponty's term) such that these can be adaptively applied to different situations. The accumulation of these "sedimented principles" leads to "know-how";[7] that is, repeated exposure to different "situations" creates a repertoire of responses that allows experienced designers to respond to ambiguities across situations. Another way of explaining this involves the contrast between "divergent" and "convergent" thinking. In the latter, design concepts are narrowed (converge upon) a promising set of solutions. In the former, design concepts spread out (diverge) as far as possible to encompass many alternatives.

"Divergent thinking" studies, such as the Alternative Uses Test or the Remote Associates Test, take inspiration from "synectics"[8] in which designers are encouraged to "make the familiar strange, and the strange familiar" or Koestler's suggestion that creativity involves the "bisociation of two mutually incompatible contexts."[9] While such approaches look as if they relate to the specific abilities of individuals, in terms of their potential to be "creative," they lack validity.[10]

A second approach to studying creativity explores how body posture or movement can constrain or influence the approach taken to the tasks. Here, body posture and movement provide "minimal embodiment"[11] that can affect divergent thinking. The view is minimal because we are considering only one aspect of the human-artifact-environment system. If we are to take this system seriously, then we need to better understand how the various elements interact with each other. Without consideration of these interactions, any account of creativity will be as limited as the information-processing or the body-based accounts. This is why it is important to understand creative practice in situ.

Hence, a third approach is to study creative practitioners in their workplace. This is my preferred approach and has been employed by researchers across a variety of domains.[12] For now, I want to look more closely at the role that problem-solving plays in discussions of creativity and more broadly of theories of cognition. The reason for this is that, in an information-processing approach, problem-solving is the sine qua non of symbolic manipulation through which a set of features needs to be internalized in order for a solution to "pop out" and reported. For embodied cognition, in the absence of mental models, how are problems solved?

Problem-Solving

Problem-solving has been proposed as a basis for explanations of creativity. Much of the work on problem-solving manipulates information-as-content. That is, people are presented with situations in which either they do not have a strategy (i.e., a familiar pattern of activity) or the strategies that they apply do not lead to a successful outcome. From this, a "problem" is a situation in which you do not have a familiar strategy for producing a defined outcome. This might be due to the features in the situation being unfamiliar to you or to some set of constraints that prevent certain actions (i.e., the "rules" by which the problem is permitted to be solved). Even before considering the features in the situation, however, it is equally important to recognize the need for an "outcome" and the actions that can be made. An incomplete Sudoku puzzle is, for someone who has no interest in Sudoku, not a "problem" but merely a partially filled grid of numbers. In the domain of artificial intelligence, Boden's[13] account of creativity has long held sway. In this, problem spaces are mapped, explored, and transformed to create new concepts, typically using strategies that have been described in problem-solving studies.

A common strategy in the problem-solving literature is "means-ends analysis." Here, the problem is presented in an initial state and the problem solver is asked to produce a goal state (end) by discovering the steps (means) to make transitions from initial to goal state. Let's take a simple example:

$$\frac{xii}{vi} = ?$$

The first thing to do is make sense of this as a problem, or to define the initial state. In terms of content, you need to interpret the symbols and know that the *vi* and *xii* stand for numbers (in Roman numerals). To define the context, you need to recognize this as something to do with arithmetic (there is an "=" sign, and the horizontal bar indicates a division sum). If one or more of these features does not make sense, then the "problem" itself is insoluble, and the solver is forced to take further steps until there is an understanding of the initial state. From the initial state, you define the goal state—here, solve a division sum to replace the "?" with a Roman numeral. Following this, the "means" are defined in terms of converting from the number system that is used here to another that is more familiar, performing some calculation, and producing a solution.

What might not have been immediately obvious from this example is that having the problem printed in front of you is a great help in attempting the solution. This reiterates a point made in chapter 1, in relation to solving simultaneous equations, which is that problem-solving experiments often make use of "external representations" but rarely consider how participants interact with these in solving the problem. It is as if these experiments assume that problem solvers do all the manipulation in their heads and then report the result. A challenge for theories of problem-solving is to explain how people keep track of the steps and the rules as they solve the problem. If all the information used in solving a problem is kept in the person's head, then one can see how this can quickly become overwhelming for all but the most practiced of problem solvers. Try verbally presenting the problem to someone and have them solve it: chances are that the statement of "*x*-*i*-*i*-over-*v*-*i*-equals what?" might need to be repeated a couple of times before they can begin to attempt this. Thus, the visual presentation of problem spaces has a bearing on how strategies can be applied.

Once you have developed an approach to solving a problem, you can apply this to similar instances. This reduces the need to "solve" the problem in future (as long as you can recall the approach). Suppose the problem is now

$$\frac{xxx}{x} = ?$$

Repeating the previous approach (translate from Roman to Arabic numerals and perform the calculation) produces an answer. But it would be much simpler to recognize that there are three lots of x above the line and one x below the line—so you can "see" that the answer is 3 (or iii) without the need for intervening steps. Information-as-context frames the problem and minimizes the need for translation. Getting stuck on a single approach to solving a problem (when there are more efficient alternative approaches) is called "functional fixedness."

Design Thinking

Design thinking is concerned with breaking free of "functional fixedness" by which particular problem-solving strategies become ossified and inflexible. Consequently, approaches to design thinking emphasize the need to continually question and challenge both the presentation of the problem and the consequences that might arise from proposed solutions.

> Design thinking relies on our ability to be intuitive, to recognize patterns, to construct ideas that have emotional meaning as well as functionality, to express ourselves in media other than words or symbols.[14]

In his "Science of Design," Simon[15] set out ideas that informed the concept of "design thinking" (although it is fair to say that design thinking is not a single school of thought so much as a loose collection of methods and manifestos that take the term in different directions). Simon's perspective echoed that of other early champions of design thinking in viewing design as an activity that could employ principles and concepts from information-processing. I am not going to review the various methods that have been advocated for these activities. Many of the methods emphasize "spiral design" processes indebted to Asimow's "iconic model,"[16] in which a vertical axis moves design from concept to prototype to product and a horizontal axis defines stages (analysis / synthesis / evaluation / communication). I am, however, interested in the ways in which design thinking has become untethered from its original information-processing moorings to become something that has more affinity with phenomenology (in terms of "user experience"). We noted earlier that Rowe's account was based on a blend of information-processing and phenomenology. However, I felt that Rowe's account has a disconnect between the practice of doing design and the "theory" illustrated by the language of problem-solving, problem statements, and "thinking outside the box." Indeed, there have been calls for design thinking to define itself in terms of "situated, embodied material practices."[17]

Design and Cognition as Multi-Objective Satisfaction

Some notions of design thinking can be traced to Campbell's[18] "Darwinian" theory of creativity. For Campbell, creativity is a form of trial and error in concept generation (or "blind variation and selection retention"). There is the implication that generating ideas "unrelated to the solution" (as Campbell advocated) would be a random process, or at least unstructured and opportunistic. Relating blind selection to problem spaces, might, if it were unchecked, lead to a combinatorial explosion which creates the sort of "wicked" problem mentioned previously.[19] Consequently, the challenge is to generate many ideas while also battling with the constraints that one might apply to make the problem space manageable. In other words, the problem space could be defined in terms of its degrees of freedom, DoF (as discussed in chapter 1). Applying DoF to design, we might define the problem space in terms of many objectives that create many constraints. One way of understanding the problem space is to recognize that constraints provide a local impetus to a global strategy.

When problems are well defined (table 2.1), producing novel solutions is less important than making sense of the constraints. Indeed, a well-defined problem most probably has a strategy that can be applied to solve it. This is why Herbert Simon proposed that novel solutions can occur only in ill-structured or wicked problem spaces.[20]

Table 2.1
Defining types of problems

Well-Defined Problems	Wicked Problems
Specific goals	Vaguely stated goals
Clear and predictable solutions	No unambiguously right or wrong answers
Clearly defined means (paths to solution)	Unstated or assumed problem constraints
Most information that is required will be available from the problem space	Require a large database of relevant information that is often difficult to access

One approach, from computer science and engineering, that has been applied to design, is the use of multi-objective problem analysis. In mathematics, multi-objective problems can be described in terms of optimization. The aim is to produce a solution that can be mathematically proven to be the best. However, it can be difficult to satisfy all the objectives. So, we select one or two objectives and treat the others as constraints (i.e., by setting limits on the extent to which the other objectives can vary). We want to optimize our objectives (maximize or minimize, depending on the outcome) by performing an action on them (i.e., in mathematical terms, we are optimizing a set of functions). In this case, we are seeking a solution that maximizes both sets of values, while also treating any other factors as limits or constraints on the solution space. This can be expressed as

$$\mathit{Maximize}: f_i(x) \quad \mathit{subject\ to}\ f_i \leq \varepsilon_i\ (i = 1, 2, q-1 \cdot q-2, q+1 \ldots, n).$$

Here, we are taking a mathematical problem and dividing it into something tractable to define the set of constraints. This process, in turn, defines a set of solutions that can be represented graphically as a Pareto Front, for pairs of objectives. In figure 2.1, each dot describes a space of solutions to the problem, using different values of two objectives, *a* and *b*. The boundary (shown as the dark dots) of this space is the Pareto Front; adjusting the constraints shifts this front and changes the solution space.

I am not going to elaborate on the mathematics here and am using the concept of multi-objective satisfaction as an analogy based on the simple observation that designers select objectives to optimize. The objectives to optimize might be defined by a collection of features in the situation, where all other features are assumed to be constrained (or held constant).

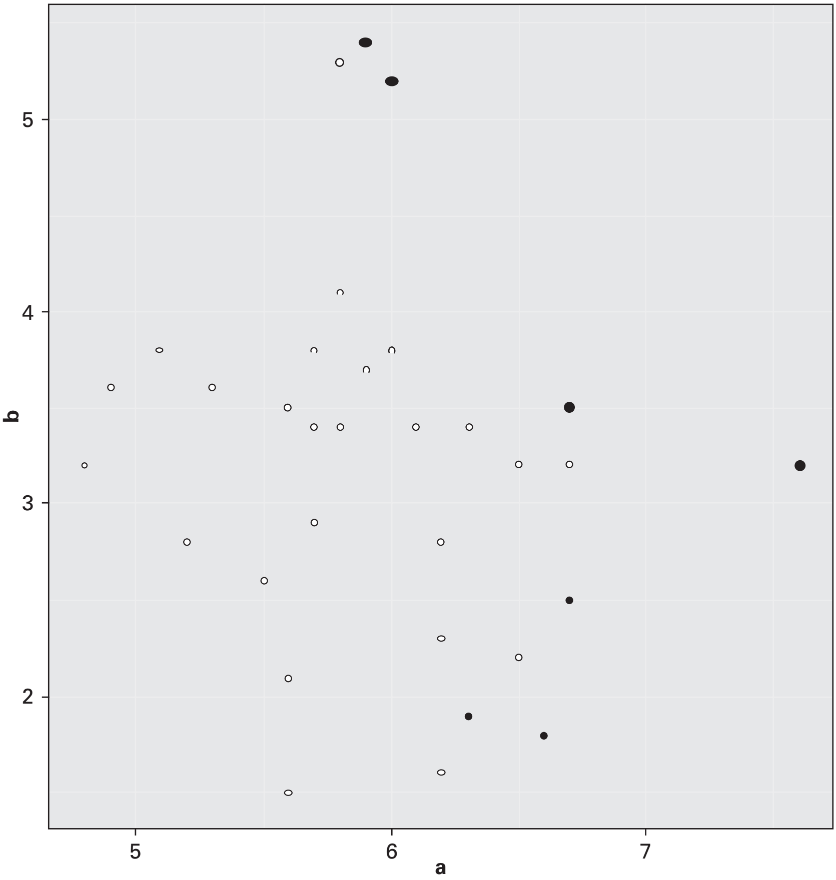

Figure 2.1
Pareto Front maximizing values for objectives *a* and *b*.

Designers are unlikely to address all the objectives in their activity, and most of the objectives will be treated as constraints on their activity (with the implication that picking the "wrong" constraints could result in poor, unoriginal, or incomplete designs). I should also note that, rather than seeking a mathematically optimal solution, designers are likely to seek satisfactory solutions—that is, solutions that involve "satisficing"[21] in terms of the objective and the constraints. In this way, we have a description of what might look like "trial-and-error" exploration during design. Design is a cyclical response to changing situational cues: "All creation . . . has the same foundation: gradual steps where a problem leads to a solution that leads to a problem."[22] In the information-processing approach, these steps involve translations into (and out of) mental models. In embodied cognition, these steps involve physical manipulations to (re)shape the problem-space.

Reitman[23] and Stokes[24] show that problem solvers can be more creative and efficient when given constraints that allow them to structure the problem space than when they have no constraint. Reitman speaks of the problem solver both identifying and breaking constraints in order to progress to a solution. However, this is not simply about violating rules. "The greater the number and complexity of the violations . . . the more the problem solver risks introducing complications."[25] The argument is that the transformation of a problem is often something that can be performed physically. Reitman characterizes this process of transformation as one of "exposition plus development plus conclusion."

For Stokes, the selection of constraints to attend to or ignore is a purposive, deliberate act. The question is how this sense of purpose or deliberation can be reflected by embodied cognition. Stokes acknowledges the possibility of spontaneous creativity, in the form of the "skilled execution" of, say, a jazz musician improvising. But even here, the introduction of constraints can lead to greater spontaneity and innovation. Designers have a repertoire of well-learned, manipulative techniques that come from their work practice. So, for example, one would expect that when confronted with materials and tools that are familiar, the creative person would respond to them in ways that exploit these well-practiced techniques.[26] Applying this to problem-solving experiments, one would expect that when confronted with physical representations of problems, people would rely heavily on their prior experience of the physical properties of objects and how they are used to find a solution. This means that the application of prior experience, in the form of repertoires of movement, becomes integral for solving problems.

The Roles of Physical Action in Problem-Solving

In the popular "Tetris" game (figure 2.2) shapes drop down a computer screen and are manipulated (rotated or moved sideways) to align with spaces in the lower layer.[27] When people play this game, they manipulate the shapes as a way of trying out different solutions. This is what Kirsh and Maglio term "epistemic action,"[28] which involves manipulating the problem space

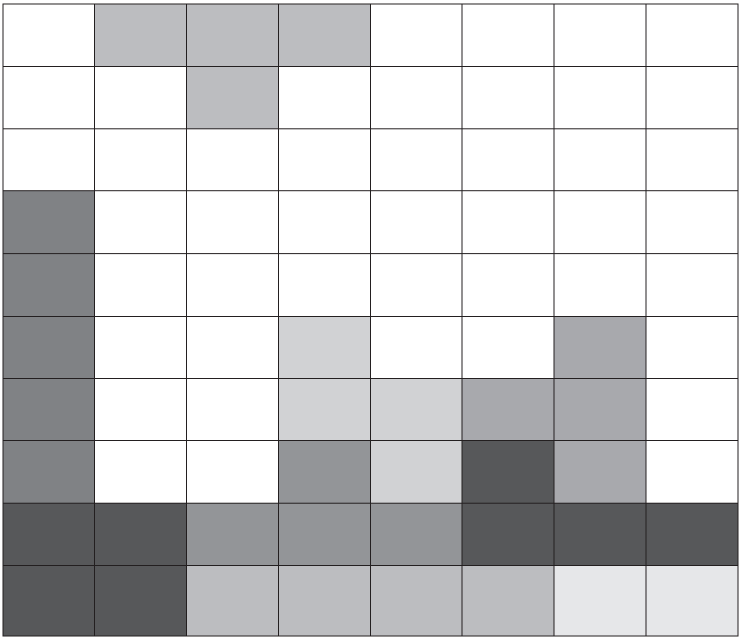

Figure 2.2

Tetris. Paul Maglio and David Kirsh show that "in Tetris—a real-time interactive video game—certain cognitive and perceptual problems are more quickly, easily, and reliably solved by performing actions in the world rather than by performing computational actions in the head alone."

in order to make available information that might otherwise be hidden. So, rotating a Tetris block helps to determine whether or not it *might* fit into an available space. Similarly, people might write (in pencil) numbers that could possibly fit a cell in a Sudoku grid before committing to a specific number. In this case, the Sudoku puzzle is (at least in its initial stages) a multi-objective problem, and the puzzler is seeking to define constraints by focusing on a single objective (or seeking to eliminate alternative objectives). Epistemic actions differ from pragmatic actions, which are the means taken to move toward a goal, by allowing people to explore the constraints in the problem space. Much of the literature on problem-solving concentrates on pragmatic actions and often disregards or trivializes epistemic actions because these are not seen as goal-directed. In the means-ends notion of problem-solving, an action that is clearly not a means to the end can be dismissed as an error (moving to a different end state) or "toying" (not moving toward any end state).

A traditional explanation of problem-solving as occurring “in the head” involves “insight.” A problem solver reaches an impasse and is unable to proceed to a solution but, after a period of time, is struck by an “a-ha” moment in which the solution “presents itself.” Often a period of time is spent away from the problem itself, perhaps day-dreaming or thinking about something else. Central to this notion is the idea that the solution appears ready formed and shifts the problem solver from the impasse to the solution. A classic experiment on “insight” is shown in figure 2.3. This involves two pieces of string suspended from the ceiling,[29] with the goal being to tie them together. If you stand between the two pieces of string, they are placed just far enough part for you to reach one but not the other (even if, with one piece held in your left hand, you step toward and reach out to the other piece). Notice that in figure 2.3, in addition to the suspended strings, there is also a collection of objects on the floor near a chair; some of these are important for solving the problem.

Typically, during the experiment the participant attempts a variety of ways of reaching the two pieces of string. About one-third of the people in the original study were able to find a solution on their first attempt. If they remained unsuccessful, the experimenter provided a “hint” by walking into the room and brushing past one of the pieces of string, setting it gently swaying. Often this was sufficient for the participant to “see” the string as a pendulum and, using one of the objects artfully placed around the room,

Figure 2.3
Maier’s two-string problem.

could swing the string so that it could be caught when standing closer to the other piece; a further third of participants solved the problem following the hint. Interestingly, fourteen people did not solve the problem at all. Furthermore, many of the people who received the hint claimed not to have noticed it. This led Maier to conclude that "the perception of the solution of a problem is like the perceiving of a hidden figure in a puzzle-picture. In both cases, (a) the perception is sudden; (b) there is no conscious intermediate stage; and (c) the relationships of the elements in the final perceptions are different from those which preceded, i.e., changes in meaning are involved."[30] The question is what does "changes in meaning" mean, and why was there no "conscious intermediate stage"?

Maier suggested that the solution required the meaning of the problem to be changed. Implicit in this suggestion is that "meaning" must involve deliberative sense-making, which is why the idea that there is "no conscious intermediate stage" was so provocative. Indeed, Maier's idea of making the string swing to provide a hint to the participant was intended to enable such a restructuring. Key to this explanation is the idea that the problem becomes restructured.[31]

One explanation for this restructuring involves "spontaneous transfer," in which prior experience is recalled and used to help solve a problem (often without conscious awareness). In this case, the problem could be solved by analogy with related solutions and actions. If this was the case, then thinking about the objects (string, pendulum, weight of objects) could be beneficial. When explicitly instructed to think about associations between objects, people tend to be better at solving the problem;[32] but asking people to define alternative uses of objects (which could include using pliers as a weight), has no more advantage to simply being presented with the two-string problem.[33] As noted at the start of this chapter, the "alternative uses" task is a common way in which "creativity" is defined, so it is interesting to note that for Maier's two-string problem it offers little benefit. For me, this points to a possible explanation of what is happening in this situation. Alternative uses require people to focus on *general* object features, which would be instances of information-as-content (in terms of the object's forms). Such a focus might not be relevant to the task at hand. In contrast, focusing on associations between objects causes them to focus on *specific* features and draws attention to information-as-context (in terms of relations between object), which could help constrain the problem space.

Clearly, solving this “problem” of attaching two pieces of string involves several physical acts and the appreciation of the behavior of objects in the world. Embodied cognition emphasizes how understanding the physical activity of objects in the world in response to our actions contributes to how we make sense of problems. Making sense of the two-string problem requires appreciation of how the physical objects interact with each other. Presenting this problem in terms of *real* objects that can be physically interacted with (rather than as a picture, such as figure 2.3) creates a different sense of the problem and the potential actions that can be performed. This means that the idea of there being “unconscious” actions makes sense only if one excludes the idea that physical action can be form of “thinking.”

The two-string problem can be defined in terms of three aspects:[34] specific features of the objects, particularly in terms of what they can be used for; combinations of these features that can be related to problem spaces; and combinations of features that permit action in these problem spaces (either cognitive or physical). For the two-string problem, the action that is required involves tying two pieces of string together, and the problem space involves bringing the two pieces close enough to allow tying. In order to produce the correct solution, people need to appreciate how moving the end of one piece of string (treating the string as a pendulum) can bring the two strings closer. Each of these aspects of representation could conceivably be performed by creating a mental model, but each is much easier to work with in terms of the objects in the environment around you. The challenge then is less about how to “think” of the objects and their relations and more about how to “see” the relations offered by these objects and how they behave.

An elegant illustration of this distinction between ‘thinking’ and ‘seeing’ was shown in variations of the Tower of Hanoi problem.[35] The basic premise of this problem is shown in figure 2.4: there are three vertical poles on which discs of different diameter can fit, and the goal is to move three discs from one pole to another, while obeying two rules: (1) only one disc can be moved at a time, and (2) a small disc cannot sit on top of a large disc.

Solving the Tower of Hanoi puzzle involves the manipulation of objects (discs on pegs) in accordance with rules. Zhang and Norman showed how changes to the problem representation can help people apply the rules.

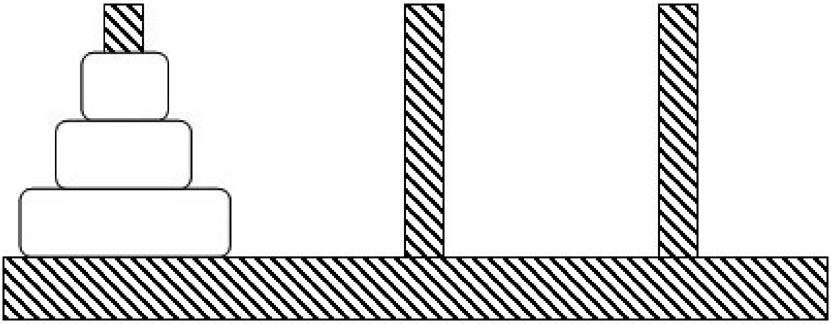

Figure 2.4
Tower of Hanoi problem.

In one version of the problem, discs and pegs were replaced with plastic oranges (small, medium, large) on plates. In other words, while the objects differed from those shown in figure 2.4, the rules were isomorphic. In another version, the problem was to move full coffee cups. The conclusion is that the more the rules are directly represented in the appearance and relation of the objects, the easier people found the problem to solve.[36] One reason why this conclusion is interesting (rather than obvious) is that the information-processing perspective would assume that objects in the world are converted into some "internal" representation in the brain. The implication from this would be that the most efficient process would be to translate the problem into a code that would apply irrespective of the presentation of the problem. After all, if the brain was going to construct mental models of a problem, why would it seek to create different models for the same problem when the only differences lie in the manner of presentation?

One way of reconciling the "internal" versus "external" debate is to suggest that the "external" representations provide memory aids, meaning that the person does not need to keep track of all of the elements of the problem in their head but rather can refer to the state of the elements in the environment. In this sense, the "external" presentation of the problem allows the problem solver to off-load some of the processing demands. So, rather than imagining what the problem space would look like if a specific disc was moved, one could make the movement and see the result. This changes the nature of the problem-solving task from one that is performed "in the head" to one that is performed "in the world." Given these findings, it is not surprising that so much of our everyday problem-solving involves some form of external representation, such as jotting down notes, sketching solutions, building simple models, or using physical objects to stand for elements of the problem. For this book, the question is whether the externalization of problem-solving is solely about supporting memory and off-loading some aspects of cognition or whether working with external representation *is* cognition.

Vallee-Tourangeau and colleagues have performed experiments in which participants solve problems either by using physical objects or with pen and paper (or tablet). In a "wolves and chicken" problem (moving pairs on animals from one side of the river to the other using a boat and ensuring that the chickens always outnumber the wolves), participants were more efficient (made fewer "illegal" moves and had lower decision latency) when they could manipulate physical objects.[37] In the "seventeen-animal problem"[38] (place seventeen animals into four enclosures such that there is an odd number in each enclosure), none of the twenty-four people who used pen and tablet was able to produce a solution, while ten of the twenty-three using physical objects produced reasonable solutions. Hint, the solution is to overlap the enclosures (so that some animals are placed in the intersection between enclosures). People in the pen and tablet condition interpreted this as an arithmetical challenge of dividing 17 by 4, and then applied the constraint that each resulting number had to be odd. In contrast, people using the models first built four enclosures (by bending pipe-cleaners into loops) and then placed model animals in the enclosures. The opportunity to work with physical objects changed the way in which the problem was conceptualized from being one about putting objects into containers rather than one about dividing numbers. From subsequent, fine-grained analysis of video recording of participants, it was apparent that rearranging physical objects often led to the opportunistic discovery of a path to the solution.[39] In this case, the physical action was not simply an aid to cognition but became a way of framing the problem, and from the framing of the problem, the solution was much easier to see. But for this to occur, there is a need to appreciate what actions are possible with the objects that are available.

Physical manipulation is important to problem-solving: as one manipulates objects (or makes sketches of ideas), so the problem space breaks down. This would mean that these physical manipulations could be an intuitive response to the problem space, thereby making it *less* structured. However, problem-solving is about finding structure, which means that, having broken down the "given" structure, the problem solver will look for ways to restructure the problem space. What does this tell us about design practice?

The Craft of Design

In a classic text on design from the 1960s, Archer[40] was concerned that "design" tended to be dismissed as a "mere craft-based skill." Like other design theorists of the time, he felt that there was a need to demonstrate that design could be described in rigorous scientific terms. The use of the word "mere" in his description is telling, in that it is symptomatic of a viewpoint that places "thought" above "action," assuming that these are separable modes of working. As should be clear from chapter 1, the position taken in this book is that acting *is* a form of thinking. In order to appreciate the antangonism against craft-based skill, we could do no better than to look at Christopher Jones's' reading of George Sturt's lovely 1923 book, *The Wheelwright's Shop*. From Sturt's accounts, Jones concludes that "craftsmen do not, and often cannot, draw their works and neither can they give adequate reasons for the decisions they take. The form of a craft product is modified by countless failures and successes in a process of trial-and-error over many centuries."[41] The implication from this is that craft moves along on an unthinking, slow, stumbling path and that it takes ages to produce a design, with this design often arising as much by accident as intention. This is an odd conclusion to draw, particularly as Sturt was at pains to point out how the design of something as humble as a wagon wheel was contingent on the wheelwright's knowledge of the type of wood that was locally available, the type of terrain that the wheels needed to cope with, the loads that the wagons would carry, and the knowledge of tradition and practice that was held in the wheelwright's shop. They might not have been able to fully explain the "tacit knowledge"[42] that saturated their craft (such is the nature of implicit knowledge and "automaticity" of skill), but this need not mean that "mere craft-based skill" is unthinking. An alternative perspective is to see the "trial and error" as the response to different sets of constraints which shifts the Pareto Front (figure 2.2) to optimize different designs.

What Archer and subsequent design theorists sought was a scientific basis for design. For Archer, this meant incorporating "knowledge of ergonomics, cybernetics, marketing and management science into design thinking," and for subsequent theorists it has meant a focus on cognitive sciences. A consequence of this, as Penny points out, is that "artisanal . . . practices have occupied a marginal place in cognitive science, because the tight and ongoing intercourse with materiality confounds notions of cognition understood as abstract reasoning."[43] The point that I would make is that artisanal (craft) practices also seem to have a marginal place in theories of design. In support of this claim, I very much like Glenn Adamson's lovely phrase "Craft only exists in motion."[44] This captures the "doingness" of craft and, for me, this translates fully into the realm of design (in all its definitions and permutations). So, the focus of this chapter is simply to ask how can we coherently capture the doingness of design? This is a simple question that is often swept aside—unthinkingly dismissed as "tacit knowledge" or, worse, seen as some form of mystical communion between designer and artifact. For me, answering this question involves elucidating a theory of "technical reasoning" that can stand against the "abstract reasoning" that worries Penny.

In work on jewelry making, my colleagues and I explored the role of symmetry in positioning stones on a broach.[45] Symmetry does not mean strict adherence. Indeed, a slight imbalance might be more aesthetically pleasing or more "honest" or more indicative of the brooch being handmade. There is, though, an interaction between the person laying out "by eye" the elements of the design and the physical interactions involved in moving the pieces until the layout is reasonable. Similar exploratory activities have been observed in the practice of architects. Here is an account of one of the architects that Reitveld and Brouwers studied:

> When RR moves the cardboard model around on the table, he lets go of the model when he seems satisfied with its position and immediately starts looking for the best position in relation to the model by moving his chair around and bending forward.[46]

From an information-processing perspective, these physical activities are of little consequence and difficult to account for. From an embodied perspective, such activities could readily be considered in terms of Kirsh's notion of epistemic action; the movement of the architect and the model provides ways to manipulate the problem space (and offers opportunities to shift emphasis between different objectives).

One of the more striking aspects of reading Rowe's *Design Thinking* is the collection of sketches produced by each of the designers. While he emphasizes the back and forth movement between convergent and divergent thinking, the sketches can be read both as "design" and as "thinking." They provide ways of representing combinations of features, eliciting the organizing principles for the designs, and allowing the designers to explore the ways in which these principles can be realized. Indeed, the question of *how* sketching is used in design is a topic of continued research activity.[47] For example, Goldschmidt discusses the "dialectics of sketching,"[48] which shifts between "seeing as" (using sketch to visualize metaphoric relations that relate to the situation) and "seeing that" (exploring the meaning or interpretation the metaphor to the design problem). This idea of contrasting a set of features with a set from another problem space calls to mind another theme running through Rowe's discussions, which has to do with the ways in which analogy is used by architects and urban planners. In formal terms, such analogies lend themselves to "pattern languages,"[49] which provide sets of features that support navigation of problem spaces. In this way, the pattern language represents (externalizes) possible points of similarity between one domain and another. In a looser sense, the use of a "mood board" allows the designer to bring "incompatible contexts" together, as, for example, in collecting seashells and using their shapes to suggest forms that can be modified for the shape of automobiles. From this perspective, sketching is a means of informally creating a personal pattern language, in which the designer works through forms and relations to help constrain the problem space or clarify the situation. In our studies of jewelry making,[50] we argued that sketches instantiate events (where an "event" is a change in the layout of affordances[51]) and that similar instantiations of events occur as the jeweler moves pieces to test different configurations, turns a stone to catch the light, or heats metal to change its color. In each of these activities, the jeweler is seeking opportunities for action in a space of constraints. This echoes Merleau-Ponty's notion of "absorbed skillful coping" or Dreyfus's "optimal body-environment relations." A more recent account of creative practice which aligns with the arguments in this book can be found in the skilled-intentionality framework, which views creative practice as "skilled engagement with affordances by the sociomaterial environment in the context of the human ecological niche."[52] We return to this framework in chapter 4 when we discuss "affordances" in more detail.

From the perspective of radical embodied cognitive science, cognition arises from the dynamic interplay between person and objects in an environment. Accordingly, sketching can be interpreted not simply in terms of the physical action of making marks on paper, but as thinking, in the form of creative and cognitive activity. This means that rather than being the result of thinking, or even an aid to thinking, the production of the sketch *is* thinking. In studies of jewelers, I have noted that sketches, when they are used, tend to rough approximations rather than fully dimensioned engineering drawings; the sketches provide an opportunity to "think through" technical problems or to communicate.[53] Similarly, sketches are "ideation drawings": "By drawing, the designer expands the problem space of the project task, to the extent of including and even discovering, new aspects, which he/she considers relevant, as much as through a subsequent interpretation of the graphic representations."[54] In addition to sketching, jewelers might lay pieces out on the workbench, experimenting with different arrangements, or might respond to fundamental aspects of the arrangement of pieces, such as their symmetry.[55] As with sketching, these physical actions can be considered as forms of epistemic action and as a way of exploring the problem space.

Fundamentally, and importantly, the "creative" act cannot be separated from the "physical" act.[56] Creative work proceeds through episodes, to use Rowe's term, in which action alternates with interpretation. For Schön,[57] design does not operate through problem-solving in the way that the information-processing approaches of, say, Simon do. Rather, "Once we put aside the model . . . which leads us to think of intelligent practice as an application of knowledge to instrumental decisions, there is nothing strange about the idea that a kind of knowing is inherent in intelligent action. . . . There is nothing in common sense to make us say that the know-how consists in rules or plans which we entertain in the mind prior to action."[58] More specifically, Schön contrasts a problem as "given" (which is typically what happens in experimental studies of problem-solving) with "problem setting" (which involves processes through which what the problem is seeking to address), the definition of an acceptable goal, the actions available to us, and so on, all of which are, in experimental studies, removed from the problem-solving or, at least, bundled together under the rules or constraints by which activity is performed). It is this "problem setting" that leads to the notion of reflective practice. Dewey (writing 50 years before Schön) made a similar point when noting the importance of a "reflective conversation with the situation."[59] We can see how the reflective conversations relates to Rowe's talk of situations and also understand how each design activity is primarily addressing a unique task in which the situation creates opportunities and challenges for the designer.

In his analysis, Schön makes use of protocol analysis, in which conversations between designers (or design tutor and student) are reported (although, ironically, his analysis involves less reflection on the part of participants than Schön's interpretation of their conversations and activity). In these analyses, "drawing and talking" are treated as parallel means through which the conversations unfold and through which designers frame and reframe the problem at hand. The conversations tend to focus on how the objectives are defined and applied in order to explore different dependencies between these. Thus, "designers might differ, for example with respect to the priorities they assign to design domains at various stage in the process."[60] Sketches and models facilitate the conversation between designer and project. Thus, the purpose of the sketch or the model is not to make visible an idea that the designer has formed already. Rather, it is the physical instantiation of constraints as they apply to the current version of the problem being addressed. That is, these activities are as much a matter of problem setting as solution presentation. From this we would expect that recognizing "constraint" and editing or backtracking within the problem space relate to an appreciation for the materials being worked with and the developing form of the design.[61] In the next chapter, I explore the "environment" as a problem space.

3 Understanding Task Ecologies

Introduction

If we accept that acting on the world *is* cognition, then this ought to influence not only *how* we design but also *what* we design—or rather how we expect designed objects to be used. The dichotomy of form and function takes us only so far—and this is because the "interaction" between person and artifact does not occur in a vacuum but within an ecology shaped by the physical and the social environments in which artifacts are used. In chapter 1, the notion of a human-artifact-environment system was introduced. Much of the subsequent discussion focused on the first two elements, and I did not say much about environment and how this influences activity. Given that developments of Gibson's original theories have tended to come under the heading of ecological psychology (an academic discipline with its own society, journals, and conferences), "ecology" plays a key role in defining this field of enquiry. But this is jumping from the word environment to the word ecology without defining either.

For Gibson, "information" resolves uncertainty over which action to perform, and this makes sense only within an environment. Indeed, it is the role of the environment in his notion of "information" that separates his notion of uncertainty reduction from Shannon's. For Gibson, the environment is not an unstructured, amorphous mass of confounding and confusing features, but a collection of resources that are salient to action. In a very real sense, the environment (from Gibson's perspective) is analogous to the abstract notion of problem space used in chapter 2. For information-as-content, "salience" is a matter of assigning meaning to the symbols that the brain creates from features in the environment, which, in turn, requires storing symbol-meaning associations (hence, declarative knowledge in long-term memory). For information-as-context, “salience” relates to the mapping of environmental features to action (which makes no demand on symbol production or storage). At this point, the reader might cry “foul” and see this as a sneaky way of using representation to support cognition—one in which the mapping of features to actions is a form of representation. However, the point is that this “representation” is *not* a symbolic reconstruction of the environment (i.e., a “mental model”) but a responsiveness to features in the environment that relate to action. To appreciate why this is the case, we need to define “ecology.”

The Environment as an Ecology

Gibson distinguishes between the environment as the habitat in which a species lives and the “ecological niche” that supports the way of life of that species. In other words, the ecological niche of a species “refers more to how an animal lives than to where it lives.”[1] So, two species of animal might occupy the same habitat, but, given their differing ways of life, might respond to the features of the habitat (in terms of what it offers for food or shelter) in different ways. “The natural environment offers many ways of life, and different animals have different ways of life. The niche implies a kind of animal, and the animal implies a kind of niche. Note the complementarity of the two.”[2] The word “complementarity” is used to reflect the relationship between the animal and its environment. The implication here is that, whereas a species might respond to ecological information in one habitat, it might not respond in similar ways in another habitat. An obvious example of this concerns ways in which some (but not all) monkeys and apes use stones as tools, for example, to crack open nuts. In this usage, an ecological niche provides the resources for action. For much of the ecological psychology that builds on Gibson’s theories, the ecological niche for humans consists primarily of the physical features that support specific actions (given specific abilities). However, he noted a distinction between the physical environment, which contains “information,” and the social conventions that are superimposed on this ecological information. For the apes and monkeys mentioned above, the presence of stones or nuts only partially predict use of tools, and the culture (in terms of a tradition of tool use) of a particular troop would also contribute to the activity. Thus, the “ecological niche” consists of both physical features and cultural conventions. This is especially true for humans: "The material structure the human-environment offers should not, and cannot be ontologically separated from the social, technical and historical lives people lead."[3]

How Do Actions and Ecologies Interact?

While the proposal that the ecology consists of both physical features and social conventions has an intuitive appeal, it creates a problem in terms of the mechanisms by which actions relate to physical features. In a simple version, the physical features of the environment provide opportunities for action (so long as the animal can perform the action). This feels like behaviorism, in which the stimulus (of the physical feature) provokes a response (the action), as when Pavlov was able to condition dogs to salivate at the sound of a bell that they associated with food. Such unmediated pairing of stimulus and response misrepresents Gibson's ideas (and embodied cognition). But it also creates a problem for an approach to cognition that denies representation: If I am not saying that a feature directly causes an action, how does the relationship between feature and action arise, especially if the action will be influenced by social conventions?

As an initial response this question, I like Ingold's observation that "much if not all of what we are accustomed to call cultural variation in fact consists of variations of skills. By skills I do not mean techniques of the body, but the capabilities of action and perception of the whole organic being (indissolubly mind and body) situated in a richly structured environment."[4] Ingold used the term "co-respondence" to describe the ways in which an animal responds to its environment and the environment responds to the animal—in other words, the respondence is mutual, and so they "co" respond. This idea takes Gibson's notion of "complementarity," which is a rather passive stance, and turns it into the active interaction between animal and environment. The animal acquires the necessary skills through its ongoing relationship with the environment and processes of enculturation through which normative behaviors are learned (say, from observing its peers). Gibson called this the "education of attention,"[5] in that the ability to attend to relevant features is something that is learned and developed over time. In this way, normative actions become those that are performed by the majority of members of a given culture.

From these normative actions, the form of an artifact reflects social conventions, or majority ways of acting. However, mass production of artifacts satisfies a host of constraints in terms of materials and manufacturing processes to produce "one size fits all" designs which might impose rather than reflect normative action. It we step back to a time when artifacts were fashioned to suit local environments and local social conventions, then the argument becomes a little clearer. We noted in chapter 2 how the work of the wheelwright reflected the environment in which wheels would be used. As another example, a survey of shovel designs in 1930s Germany found over 12,000 alternatives.[6] While you might expect some variation for different tasks, and some variation for the use of different materials and cultural traditions across different regions, this is still a staggering number and suggests that the form that these shovels took was influenced by more than their basic functions, such as digging or lifting. The variations included the size and shape of the blade (reflecting the type of materials that they were required to work with) and the length of the handle (reflecting the physical activity involved in using the shovel and the normative way of performing this activity in that particular region). That is, "normativity" always relates to a concrete situation.[7] Acting normatively requires correspondence to artifacts (and other people) in *that* particular situation—for example, climbing *those* stairs in *that* building.

Responding to Ecologies

We noted previously that in order to manage the uncertainty in interactions, the elements of the human-artifact-environment system need to adapt to each other's behavior or to act is if their behavior is constrained. Taking the idea of "ecology" further, Newell[8] argues for three constraints on human behavior:

1. organism, which is defined by the size and shape of the actors, bodies and their capabilities to control their movements;
2. task, which is defined by the set of acceptable outcomes of action, in terms of how an intention is met but also in terms of how efficiently an action is performed;
3. environment, which is defined in terms of objects respond to general laws of physics, such as gravity or reflection.

From constraints 1 and 2, one can define combinations such that the perceptual capabilities of the organism can interact with the physics of the environment. The point is that embodiment is more than simply the possession of a body and has much more to do with the ongoing, reciprocal engagement between organism and environment in pursuit of task outcomes. Constraint 2 highlights that task outcomes are defined by the normative social conventions, in terms of what defines “acceptable.” You could, for instance, eat from your knife, but it would not be acceptable in polite company. Of course, there are situations in which it might be acceptable—for example, if you are slicing an apple with a knife, it might not be a problem to raise the slice of apple on the knife to your mouth. This returns us to point that normativity is always related to a social situation.

Salience, Action, and Information-as-Context

An abiding question is how can features from the environment be defined as salient? In other words, what can we say about “meaning” if we are replacing content with context in our definition of information? In Gibson’s terms, salience is directly tied to the action to perform and, crucially, does not require mediating activity (in the form of information-as-content). The broad premise of Gibson’s work is that information exists as an array in the environment and that we respond to elements in this array according to the actions that we are performing. Cognition is thus less about information processing and more about response tuning.

Gibson’s best-known account of this was in the form of optic flow experienced while moving through the environment. Gibson’s initial thoughts on optic flow were inspired by discussions with pilots and his own experiences in flying and making training films for the US Air Force. As figure 3.1 illustrates, when landing an airplane, the pilot has the experience of heading toward a fixed point and the rest of the environment seems to be moving away from that point and “flowing” around the pilot.[9]

To account for this experience, Gibson proposed that patterns of light reaching the retina constitute an optic array that contains visual information from the environment (including the position of objects). The “flow” of this optic array results in a “textured gradient” that indicates speed (of self and moving objects in the environment) and relative distance. Responding

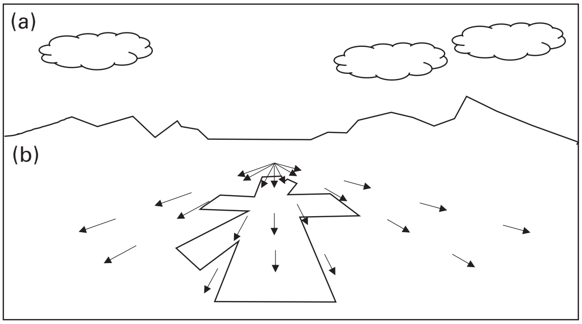

Figure 3.1
Optic array for a pilot landing an aircraft.

to changes in the optic array would be sufficient to support action, without the need for "information processing." Indeed, the very idea of optic flow is difficult to translate into information-as-content because it is not obvious which features of this flow can be represented symbolically. The labeling of features with symbols conjures up Garcia-Marquez's fictional village of Macondo where, during a sleeping sickness that robbed people of their memory, signs were hung on everything: "This is the cow. She must be milked every morning so that she will produce milk, and the milk must be boiled in order to be mixed with coffee to make coffee and milk."[10] Ironically, the sign conveys information-as-content with little clue as to how the actions (milking, boiling, mixing) ought to be performed.

The direct experience of the optic array does not require processing of information from the environment into a "mental model." Where there is uncertainty in visual information, this can be resolved by altering the experience of the optic flow, such as by moving the eyes or the head. The sampling of the optic array depends on the combination of the movement of the person and the content of the environment. Gibson proposed that people can "tune into" the optic flow in a specific environment in such a way as to allow their visual perceptual system to find resonance with the environment. From this, the environment is experienced directly, without the mediation of information processing. It is only a small step to associate elements in the optic flow with corresponding actions. For example, in the concept of "safe field of travel" (figure 3.2),[11] Gibson and Crooks suggested that the optic array could be considered in terms of regions around, for

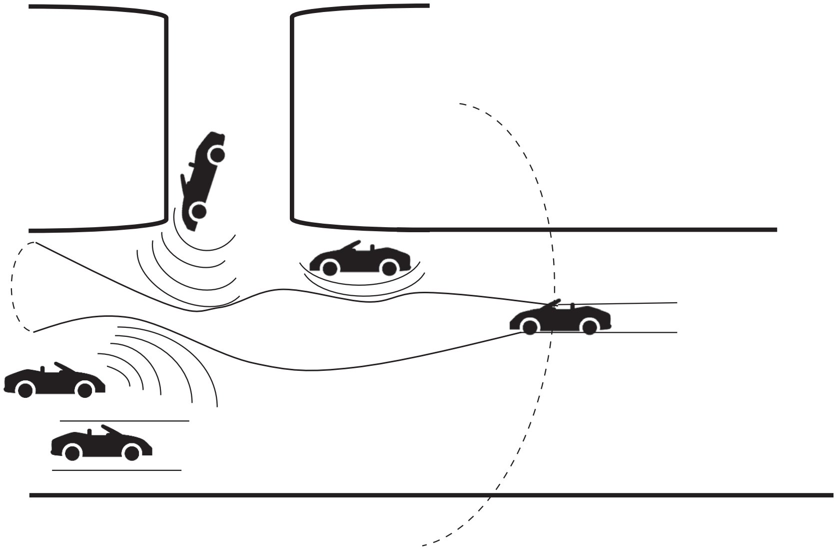

Figure 3.2
“Field of safe travel” for automobiles.

example, the automobile in which you are riding (either as driver or passenger) and those of other automobiles (or other objects on the road). Such “safe fields of travel” are defined primarily by relative position and speed in order to avoid collisions. This “safe field of travel” does not require the construction of a mental model of objects in the environment, but a continual sampling and adjustment relative to the optic array. In these examples, the environment is viewed as a complex pattern of “information” that changes in response to the actions that the person performs (or the actions of artifacts, people, and other elements). Activity in the environment is a continuous process of adapting to the available “information,” either in terms of acting to avoid collisions, or to resolve ambiguity, or to modify the environment.One can conceptualize activity in the environment as a closed-loop (cybernetic) control system. Significantly, separating the elements of the system (i.e., human-automobile-environment) into constituent components is not feasible[12] and this points to another distinction between information processing and embodiment in that the borders between elements are permeable.

Brunswik's Lens Model

While Gibson's work is concerned predominantly with the relationship between movement and salient "information" (-as-context), a parallel theory explored the role of "information" (-as-content) in decision-making. Brunswik's lens model[13] of decision-making broadly involves two elements:

1. a set of cues in the environment that can be objectively associated with a given decision outcome, and
2. the selection of those cues according to a human decision.

Accordingly, cues (i.e., features of the environment) can be assigned a diagnostic value (i.e., a correlation or weighting in terms of the relation to the "correct" decision), which is the "true cue" validity. This can be contrasted with the cues that the decision maker chooses, or the "observed cue" validity. The performance of the decision maker can be evaluated by the correlation between these two forms of cue validity. Brunswik's proposal is that the environment contains cues that can be defined in terms of their salience to making a correct decision (figure 3.3). Contrary to "optic flow," Brunswik's model requires a series of steps by which the environment is perceived (in terms of an image on the retina), and this image forms the basis for probabilistic inferring of salience. In this way, there are mediating states through which "information" is translated, and this, although not directly related to symbolic information-processing approaches is part

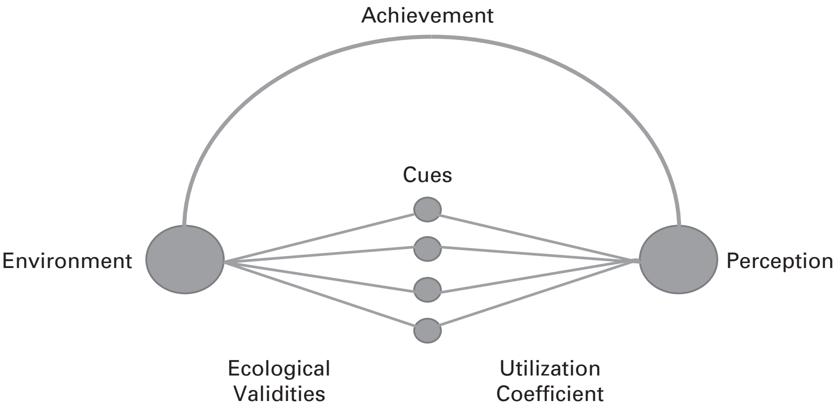

Figure 3.3
Illustration of Brunswik's lens model.

of the same family and hence, not an embodied cognition approach. However, there could be some benefit in considering how these might overlap.[14]

For Gibson, "information" arises from the one-to-one correspondence (complementarity) between an invariant feature in the environment and the action that the organism can perform (which, in turn, relates to the concept of "affordance," which will be considered in chapter 4). This means that there is no requirement for "translation" and that perception is direct. For Brunswik, the environment contains cues that are "context-free." In a study of stair-climbing, short and tall participants were asked to determine whether stairs could be climbed easily or not.[15] Plotting the data simply in terms of riser height (as a "context-free" metric) showed little correspondence across the two groups. However, plotting the data in terms of the ratio between riser height and leg length showed a neat correlation that was consistent across groups.[16] In other words, "information" is not context-free but rather involves an organism's specific frame of reference.

Both Gibson and Brunswik recognize that "perception" occurs in an animal-environment system and arises in terms of the capabilities of the animal to respond to its ecology.[17] Indeed, it is these aspects that influence Brunswik's emphasis on "ecological validity." For Brunswik, a "representative design" (of an experiment) would match the correlations between cues and outcome found in the real-world setting of the task. The reason why this is critical is that, for Brunswik, people learn to weight cues on the basis of successful outcomes (where, over time, they receive feedback on the accuracy of their decisions); indeed, this might allow the experienced person to discover relationships between cues (i.e., contiguous cues), which can allow predictions to be made accurately by sampling fewer cues (i.e., if cues are related, then one needs to sample only one rather than all of them). However, these correlations (between cues and between cues and outcomes) depend on the environment and on social convention; if the environment changes, the relations between cues and outcomes might change, with the result that performance will deteriorate; if social conventions change, then the definition of a successful outcome may change.

For Brunswik, the ability to apply strategy distinguishes experienced and inexperienced decision makers. As decision makers gain more experience of the environment, their ability to use the relationship between cues and outcome improves.[18] Indeed, the implication (in Gibsonian terms) is that the human will become tuned to specific cues (at least in familiar environments).

As they gain more exposure to these cues, combinations of cues could allow redundancy (which can improve reliability of decision outcome). If there are too many cues, then it becomes difficult to define salience. However, it is plausible to assume differential weighting of cues and that people will seek out those that are more salient to the task, and this is the basis for contemporary theories inspired by Brunswik, such as Gigerenzer's "take-the-best" heuristic.[19] In much of this work, the number of relations that are used is very small, amounting to two or three (with more than three being deemed "high").[20] Relating the number of cues to the discussion of multi-objective optimization, the finding that people tend to work with highly limited sets of relations (or cues) aligns neatly with the suggestion that design might focus on a limited number of objectives. In this, the definition of an "objective" (in terms of design) becomes a means of defining relations between cues, which implies that the weighting of cues (defining their salience) can have an impact on how these are sampled and selected.

Much of the research employing Brunswik's lens model focuses on static decision problems, rather than activity in dynamic environments. Consequently, the relationship between the "cue" (from the environment in which the decision is to be made, e.g., in terms of sets of "information-as-content" provided in an experiment) and the "perception" by which the cue's salience is defined becomes caught in a discrete, static task. For Gibson (and embodied cognition) the tasks are dynamic (like the "field-of-safe-travel" problems discussed earlier), and this makes it difficult to distinguish the cues in the ways in which Brunswik lens model assumes.

Exploring dynamic decision making, an experiment simulating an anti-air warfare coordinator,[21] performance was modeled using the lens model. The results suggested that "good" performance involved pattern recognition of salient information. Furthermore, training, feedback, and practice that allowed participants to refine the heuristics by which cue salience was defined improved performance. That is, performance did not involve appeal to mental models, but focused on perception and action. For Kirlik,[22] the mapping between the perception and environment becomes direct—that is, a one-to-one mapping—with experience. What is attractive about this is that it brings the possibility of rapprochement between the Gibsonian and Brunswikian approaches (although, of course, there remain clear and obvious differences in methodology and underlying theories) in that increasing the experience of an environment is likely to increase the consistency by which specific cues are deemed salient.

Recognition-Primed Decision-Making

One implication of the research using the Brunswik lens model is that people ought to attend to the most salient (rather than all) the available features, and that selection of features could be influenced by action (in a reciprocal manner to the action being influenced by the attended features). Recognition-primed decision-making (figure 3.4) interprets a situation in terms of "relevant cues." These are evaluated in terms of "plausible goals" and "expectancies" (the latter drawing on a mental model, or schema, that reflects experience of previous situations). In this model, when an action is defined as suitable, its outcome is simulated prior to implementation. However, work with expert decision makers tends to suggest that this simulation phase is rarely implemented.[23] This implies that decision processes (particularly in the highly dynamic, risky, and ambiguous situations that Klein and his colleagues study) is a matter of perception-action coupling. Indeed, an alternative version of this could be proposed in which there is no requirement for schema but in which the salience of cues (acquired through experience) are defined by the experts;[24] analogous to the description that the Brunswick lens model offers. Decision-making that is "automatic" and that operates through perception-action coupling not only provides an explanation of the fast "intuitive" decision-making that is required in emergency situations (and which people might find difficult to articulate because of the reliance on "tacit" knowledge) but also explains the role of expertise (gained from experience of many different types of situation) in shaping response.

As we saw in chapter 2, the way a person responds to a problem is highly dependent on the way in which the problem is presented and the physical actions that are possible in that context. In the words of Vallée-Tourangeau and colleagues, "A reasoner is embedded in a certain task environment that together configures a certain cognitive ecology within which certain cognitive abilities are manifested."[25] That is, the problem solver becomes immersed in the context of the problem, and this very immersion creates opportunities for acting and thinking. The implication is that being able to physically act in the "cognitive ecology" of the problem can be beneficial for problem-solving and,

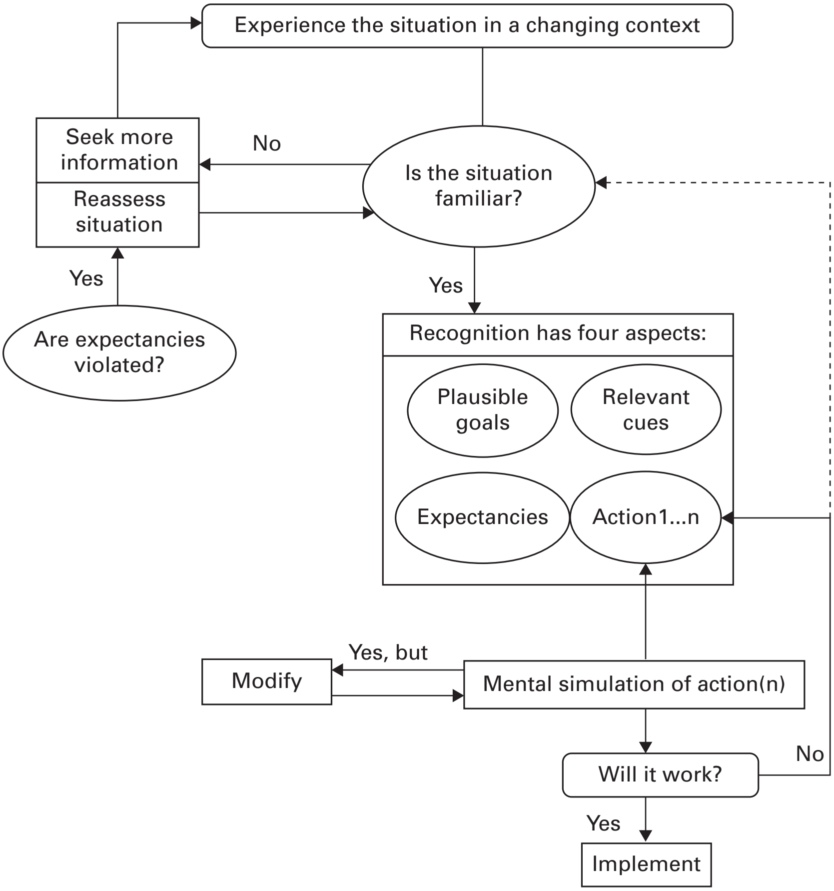

Figure 3.4
Recognition-primed decision-making.[26]

indeed, change the nature of the problem and the strategies used to solve it.[27] For now, the question to pursue concerns how "meaning" can be arrived at from an information-as-context perspective in terms of a "task ecology."

A "Task Ecology"

If radical embodied cognitive science is to be useful for design, it needs to reflect the ecology in which activity occurs and to do so in such a way as to provide rigorous, testable descriptions that could inform design decisions.

For the human-artifact-environment system, we ought to define activity not simply in terms of the single human using the single artifact, but in terms of the "context of use" (or ecology) of activity that carries a sense of "flow."[28]

From the initial discussions in this chapter, it should be clear that an environment can be described in functional terms—that is, relative to the goals and actions of a person, which constitutes the task ecology in which activity is performed. So, when driving a vehicle, the environment has features that support functions, such as the terrain (whether the ground is rough, muddy, slippery, whether there are hills, and so on), obstacles (other road users, pedestrians, objects on the road, and so on), and the path to follow (e.g., type of road, markings on the road, traffic signs, instructions from a navigation system, and so on). The salience of features is relative to the overall goal of the activity. This observation raises two key concepts. The first relates to the selection of salient features, and this could be considered in terms of their Gestalt (pattern) for the system—that is, the relationship between the whole "environment" and the "part" (pattern) that is salient to the goal. In this case, the goal becomes interpretable in terms of the system that performs the activity within the environment in which the activity is being performed. This relates directly to the idea of a human-artifact-environment system. The purpose of the system arises from a combination of the functions that it is configured to perform (some of which might interact with others) and the state of the environment in which it is operating. Each of these provides constraints on the system's operation. In this respect, the design should present only those combinations of features that are salient to the function at hand.

A second concept that arises from this notion of task ecology as a set of constraints is that the "meaning" of features (in the environment) becomes salient to activity. In this respect, the design of a user interface (see chapter 5), should present salient features in a manner that enables means-ends analysis. In this respect, a "task ecology" consists of the set of features in the environment that are relevant to a specific goal (performed by a system that is able to recognize and act upon those features), and so the design of a user interface would present salient features in a manner that allows direct perception of system state and definition of appropriate action. Here, only that part of the system relevant to the situation is presented, and this decomposition (from whole to part) becomes integral to design. The overarching framework of this notion of task ecology echoes the following quotation from Gibson:

> Things are components of other things. They would constitute a hierarchy except that this hierarchy is not categorical but full of transitions and overlaps. Hence, for the terrestrial environment, there is no special proper unit in terms of which it can be analysed for once and for all. There are no atomic units of the world considered as an environment. Instead, there are subordinate and superordinate units. The unit you choose for describing the environment depends on the level of the environment you choose to describe.[29]

The environment is a "space of possibilities," in which different constellations of features arise to afford different functions that can be realized through combinations of activity)[30]. The point is that the very of ideas of data or activity are seen, not as defining the environment, but as low-level components that constrain and support the functions of the system in that environment. So, how do we make sense of task ecologies as problem spaces.

Studying Task Ecologies

An ecology is a mesh of features defined in relation to different actions. Depending on the capabilities of the person, the experience of these relations will alter. This seems to me to be a reasonable formulation of what Merleau-Ponty called an "intentional object." which describes that collection of features in an object that correspond to a given intention for its user. I should note that Merleau-Ponty did not assume that the "intention" was some teleological impetus (that is, the "user" need not set out to achieve a definite outcome in "using" the object); rather, he saw the "intention" as becoming realized through the ensuing interaction.

In her analysis of how people interact with technology, Suchman[31] observed pairs of participants trying to use a complicated photocopier. Her study employed conversation analysis to reveal the ways in which people respond to the changes in state of the machine in relation to their own developing and emerging goals, as well as the ways that sense is made through their interactions with the device and each other. An insight from this work was that people respond to opportunities offered by the situation (or, at least, their interpretation of these opportunities) to choose which action to perform—and then, in retrospect, frame these choices as a plan. One can read the transcripts of the conversations as ways in which people struggle with the photocopier as an object that is "present-at-hand," rather than as something that facilitated their activity as "ready-to-hand."[32] Furthermore, her insistence on reflecting behavior in its situation (rather than only on the "tasks" performed by users of technology) creates a richer and more elaborate perspective to the one that Human-Computer Interaction (HCI) had typically employed before this work. However, by capturing specific instances (no matter how much detail the reports contain), ethnomethodology (in HCI and design in general) becomes "fact dependent."[33] This means that any generalizable principle (or design guideline or requirement) can be only an ad hoc response to the specific instance. So, the approaches offer no "guide to discovery" in which predictions can be made or conclusions empirically tested. In a sense, the collection of instances become like stamps to the philatelist. Such criticism can be leveled with equal force at other methods that rely on observing people as they work (and so is equally applicable to the ergonomics methods of task analysis). Addressing this criticism can involve either a sufficiently long period of immersion in the domain that is being observed or the development of a theory that can provide a "guide to discovery." In terms of the former, there is some concern with the manner in which HCI is currently practicing ethnomethodology, as expressed by two people who had been involved in the field from its early days. As Sharrock and Anderson write,

> We were part of the effort which opened up the promise of ethnographic fieldwork for software development. We still believe that it has a lot to offer. However, in the intervening years, we have watched as what can only be regarded as old-fashioned customer relationship management has masqueraded as ethnography. As the consultants have moved in, the canons of fieldwork rigour have been eroded. Now, it seems, any kind of conversation with any kind of user/manager/customer can be called ethnography.[34]

The emphasis on the richness of the context of use that Suchman called for can be seen in similar trends in philosophy, with phenomenological interviews,[35] and in the field of naturalistic decision-making, with its use of critical decision method interviews.[36] In each of these domains, the objective is not merely to elicit the sequence of tasks to perform or the information that is attended to, but to situate activity in the settings that gave rise to it. So, for example, a critical incident interview might encourage interviewees to reimagine the sounds, smells, environmental conditions, and so on in a situation before asking about how they dealt with an emergency. Similarly, Gallagher 's account of "expert performance" (in dance, sports, or music)[37] emphasizes *performative awareness*, by which the expert can attend to, and reflect upon, the movement of their body in ways that the less experienced person might not notice or appreciate. Having said that, the objective of the phenomenological interview and the critical decision method is not simply to reconstruct a specific situation but to obtain generalizable observations that can inform theory and design. That is, the aim of a phenomenological interview is not to capture idiosyncratic, personal experience, but to "capture the invariant structures of experience."[38] I wonder whether ethnomethodology can focus too much on the specifics of the situation, leaving the designer to translate the comments to a design concept. In this, the analyst becomes guilty of abrogating responsibility for translation of their findings. By "translation," I mean the abstraction of general points from specific accounts that then become formalized into design briefs, requirements, specifications, and the like. Because these approaches to analyzing technology in situ are fact-dependent, they have neither the aim nor the ability to serve as "guides to discovery." This means that someone needs to undertake the translation. Having the analyst integrated into the design team can dramatically reduce these translation problems. Certainly their presence can help in explaining the analysis, resolving confusion, and addressing contradictions. But the very act of "reducing" the analysis to a set of "user stories" or "persona" could become antithetical to the capturing of "real experience." Equally, because the analyst becomes a "proxy" for the end-users, there might be a danger that it is the experience of the analyst that is used to "stand in" for these users. This can be even more telling when the "analyst" is relying on previous experiences (as when subject-matter experts work in design teams). Continually reflecting the analysis (and its interpretation) back to the people who were studied or will be actually using the design can ease some of these problems.

Suchman's work set the tenor for a whole field of ethnomethodology in HCI. In this field, the environment is "meaningful primarily through the ways we interact with it."[39] At first glance, the emphasis here is on "the ways we interact" (and, indeed, Dourish's influential book is called *Where the Action Is*). However, ethnomethodology pays less attention to physical action (at least in terms of "action" as it could be described from fine-grained analysis of movement) and to the environment (at least in terms of the rich configuration of features that it offers) and more attention to the verbal account of the experience of action. This is not to say that movement or the environment is ignored; in many of these studies, scenes from video recordings of people interacting with technology are used as data and complement the verbal accounts. However, I have two concerns about this approach. The first is that the experience being studied comes partly from the words of the person performing the action and partly from the author's description of the video. This reiterates my point that the analysis is less about the action itself and more about the description of the action.

Aside from concerns over how "ethnography" might be practiced or how it might provide a "guide to discovery," my worry is deeper and concerns the focus on "meaning" rather than action per se. While Gibson regarded the environment in objective terms (that is he took a realist position in terms of ontology), he also wanted to be clear that the experience of the environment is shaped by individual capacities—that, the environment contains features that exist whether they are "experienced" or not. But when we consider the ways in which "experience" is collected, through ethnomethodology, as verbal reports, then "action" becomes the input to an interpreted account, rather than an activity in its own right. If we take the phenomenological position of Merleau-Ponty in as considered a manner as possible, then it would seem that asking people to provide verbal descriptions becomes almost the *least* effective way of capturing their experience. At best, such verbal descriptions are loose, informal, and incomplete accounts. The use of verbal description to derive meaning runs significant a risk of tipping from a phenomenological account of experience into one that is indistinguishable from the nominalist accounts of information processing from which it seeks to depart. This is because a verbal account, by definition, enforces a commitment to symbolic representation. So, this creates an impasse. Granted, ethnomethodology, the phenomenological interview, or the critical decision method take great pains to create a broad a view of the context. Granted, each of these approaches argues for the importance of context and (as has been argued throughout this book) acknowledges that action can only be considered in context. But there is a danger that an account that is heavily contingent on verbal descriptions (either of the participant or the observer) runs the risk of reporting the external, describable characteristics of the environment or the action and misses the underlying aspects of the context (in terms of the relations within the human-artifact-environment system). In a sense, my concern is that this approach could capture information-as-content (through its verbal description) and miss information-as-context (in the Gibsonian sense). In the words of Sharrock and Anderson quoted earlier, the risk is that ethnomethodology in HCI becomes "good old-fashioned customer relationship management."

Another way of seeing this is that the approach could reflect aspects of epistemology without fully appreciating the ontology. To do the latter would require a definition of the situation in terms of those features that are salient to the action.

The use of verbal reports as data[40] has been criticized in terms of the relation between such reports and information-processing. For instance, there is concern that people will describe only those aspects of their activity of which they are consciously aware (making it difficult to capture tacit knowledge or to fully reflect the impact of the environment on activity), or that people will describe only those aspects of activity that can be put into words (making it difficult to capture procedural knowledge), or that people might alter their activity to make it easier to describe. Such criticisms can be traced back to the challenges to introspection raised by the early experimental psychologists. In a sense, such criticisms point to a schism between those approaches that favor formally observed behavior in controlled settings and those that emphasize the importance of individual experience. This could be seen as a fight between laboratory-based approaches (central to much of the information-processing school of cognitive science) and ecological and experiential ones. But, from the perspective of embodied cognition presented in this book, the distinction is not quite as clear-cut.

Several of the studies that I have cited, in this chapter and in chapter 2, rely on conversation analysis to combine the words as they are spoken by people engaging in everyday activity. For some forms of analysis, video alone might be informative and instructive. For example, William Whyte created documentaries, such as *The Social Life of Small Urban Spaces*.[41] He provides a commentary on the activities of people, say in a public park, using different camera angles, to draw attention to aspects of the environment that we often miss or take for granted. The film allows us to "see the general pattern of behavior or sociocultural practice in this place, but when we zoom in we find a great variety of ways in which people engage with the various action possibilities the park offers."[42]

Combining video from, say, a head-mounted camera with verbal commentary (recorded either during the video capture or through interview afterward) provides a further means of capturing some aspects of the experience of the skillful coping of people as they undertake work in "real" settings as opposed to the artificial conditions of the laboratory.[43] The imagery (video or stills) gives context for the conversation analysis. In this way, the verbal reports takes precedence, with stills from the video presented to illustrate what is said. Notwithstanding the problems raised concerning verbal reports in the preceding discussion, it feels to me that relegating the video to illustrative purposes loses much of its benefit in capturing the dynamics of the interactions.

Richer analyses can be obtained by combining these words with stills from video recordings of the task performance.[44] This provides the opportunity to explore the temporal setting of the conversation as well as its physical and linguistic unfolding (figure 3.5).

From the perspective of “ecological validity” (as defined by Brunswik), the challenge is to adequately capture the cues that people use to inform their decision-making. Contemporary theorists such as Gigerenzer apply decision tasks in which cue selection is an essential feature of the experimental design. This allows consideration of the manner in which people select between the cues available to them. This could be done simply by hiding the cues until people select them—for example, by having information on a computer screen that becomes revealed only when the person clicks on it or through the use of eye-tracking. Combining these two approaches, my colleagues and I explored how people choose from a set of options in a simulated credit card fraud analysis task.[45] The layout of the information (in a grid on a computer screen) satisfies Brunswik’s notion of ecological validity (in that the aim was to present the correlations between cues and outcome found in the “real-world” task).

Eye-tracking was used to explore strategies of information sampling when people copy a pattern presented on a screen. Participants tend to look at patterns at strategic moments.[46] For example, an initial glance might be to identify the color of the block, and a subsequent glance might be to determine its precise position in the pattern. Such a “minimal memory strategy” shows how people respond to task constraints in their sampling of the environment. Thus, the manner in which eye movements are performed is influenced by the type of task[47] Applying eye-tracking to working environments is more challenging but can reveal insights into the strategies that people use to sample their environment in order to make decisions.[48] Using eye-tracking in vivo, however, can be challenging to implement,[49] not least because the human eye is in constant motion ,and this means that where it pauses (fixates) need not be the point at which “information” is obtained (one could, for example, be gazing into space rather than looking

	0	.02	.04	.06	.08	.1	.12	.14	.16	.18	.20	.22
P5 ()								[speech not audible]				
P1 action	Sitting on edge of desk 1											
P2 (blue)												
P2 action		Marks map of UK on desk 1 [slide 5}					Stands			Walks to whiteboards		Writes on whiteboard
P3 (white)								Sunny Jim?				
P3 action	Looking at statement of harbor master log and harbor								Walks to whiteboard			
P4 (grey)	Yeah		Yeah		Yeah yeah yeah				Sunny Jim is the name of the boat.			
P4 action	Looking at statement of harbor master log and harbor map										Walks to whiteboard	
P5 (red)	…this boat (0.2)…ah…that docked here…which is really close to the car park (0.1) and there's another…ah…											
P5 action	Looking at statement of harbor master log and harbor map											

Figure 3.5

Communication between participants in a sense-making exercise.

directly at something). However, like the approaches that combine conversation analysis with video recording, eye-tracking provides a means by which investigation can attend to both the properties of the environment (as these impact on the actions being performed) and the dynamics of the activity.

In a study of the micro-materiality of the handling of surgical instruments, Heath and colleagues noted how little talk occurs. Consequently, their analysis uses "a series of images accompanied by brief descriptions in relation to the timing of the particular activities. . . . In fragment 1 for example, the activities in question commence 16 seconds into the fragment and last for just over 3 seconds."[50] Regarding the activity of a scrub nurse passing a dilator to the surgeon, they note that "the scrub nurse clasps the dilator at both ends and passes it horizontally to the surgeon. The surgeon is able to grasp the dilator with his thumb and forefinger in the center of the rod. Without adjusting or repositioning his hand, he immediately inserts the head of the dilator into the patient's oesophagus."[51] This evocative description allows them to discuss not only the physical actions that are performed but also the knowledge that scrub nurse and surgeon need to have in order to collaborate, with each making sense of the other's actions within the social milieu of the operating theater. Here, the task ecology is the social and physical environment, rich in the culture of surgical practice. If we are to understand a given task ecology, then it is studies like this that will provide the level of detail required. However, while the description (of the scrub nurse and the surgeon) is detailed, it lacks the nuance and richness that could be obtained from the original video (or from data that would fully describe the performance of the activity).

The focus on "materiality" in these studies aligns neatly with the concepts in this book. In this sense, materiality concerns the questions surrounding the physical properties of artifacts and the impact of these on how they are used.[52] Recent work by Dourish[53] extends the concept of materiality to the ways in which digital technology represents and allows interpretation of information. Notwithstanding my discussion of the different uses of the terms of information (for Dourish's analysis the focus is on information-as-content), his work emphasizes the need to appreciate how interpretations as intended by the designers or managers of technology might differ from those of its users. To return to the example of handling surgical instruments describe by Heath above, each of the actors in this activity might regard the same artifact as having different meanings, which then compel different opportunities for action. The physical appearance of the artifact and the environment in which it is used define, in a sense, the ontology of the actors. Their interpretation of the artifact forms their epistemology, which reflects the ecological niches in which each of them defines the appropriateness of action. The implication is that the mere physical appearance of the artifact is not an imperative to action; one cannot simply assume that form dictates a single function. Nor, as I argue in the next chapter, does it make sense to speak of “affordance” being the property of the object. Rather, the materiality of the artifact is captured by both its form and meaning to the person using it for the purpose that person intends.

To this end, adding sensors to the person or the artifacts they use[54] allows fine-grained analysis of activity. While much of the work on human activity recognition and analysis focuses on identifying specific, discrete actions, it makes more sense to understand activity in terms of the manner in which movements balance between consistency and variability, as, for example, in terms of Bernstein’s notion of dexterity. Rather than asking “what” action is being performed, it is more useful to ask “how” actions are being performed, and this involves the application of dynamic systems approaches to the study of human activity. In such approaches, the analysis considers the temporal variation in the control and organization of patterns of movement, often using metrics related to entropy. Radical embodied cognitive science makes extensive use of these metrics and concepts, and this will be explored in chapter 7. In the next chapter, I focus on the ways in which perception-action coupling in the use of physical artifacts has been explored using the concept of affordance.

4 Affordance

Affordances Are Neither Form nor Function

Don Norman said that it is obvious what is meant by “affordance”: “Plates [on doors] are for pushing. Knobs are for turning. Slots are for inserting things into. Balls are for throwing or bouncing.”[1] What is deceptively attractive about this list, for design at least, is the implication that the physical form of an artifact dictates which action to perform. The job of the “designer,” from this perspective, is to ensure that the form of the artifact supports the desired function. But if Norman’s definition is correct, then it offers little beyond the design credo that “form follows function.” While there has been vociferous debate on the difference between form and function, the credos themselves shed no light on the practice of design, partly because speaking of “form” or “function” in the abstract makes little sense. One route out of the cul de sac of abstract function is to declare the “function” to be whatever culturally significant “meaning” is applied to a given form. In this way, however, we stumble back into information processing, where “meaning” has a symbolic representation. To preempt the argument that will be advanced in this chapter, assuming that “affordance” is a property of the artifact ignores the situation in which it is used, in terms of the capability of the user or the features of the environment. Further, the notion that there is a “desired function” that dictates how the artifact “should” be used implies that the best way to use an artifact can be dictated by the designer through the form of the artifact.

Looking at the holes in the handles of a pair of scissors, you might guess (if you had not seen such an artifact before and you are right-handed), that the smaller hole is for the thumb and the larger hole is for two fingers. Holding the scissors in this way allows movement of thumb and fingers to open and close the scissors. The form (holes in the handle) indicate function

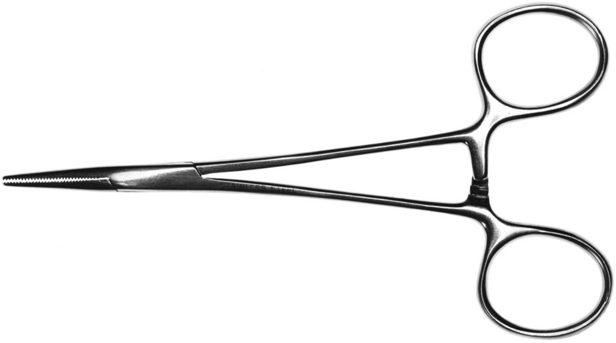

Figure 4.1
Needle holders. Needle holders do not necessarily "invite" an appropriate grasp, unless one has knowledge of the function (activity) to which they can be applied. The placement of thumb and ring finger is natural for the practiced surgeon but not necessarily intuitive for the inexperienced user. Surgical Instrument (clamp, needle holder) is licensed under CC BY 4.0.

(how to grasp the artifact and, once grasped, how this artifact might be manipulated). What happens if we apply this interpretation to a visually similar but less familiar artifact, such as the needle holders used in suturing (figure 4.1)? The holes in the handles of needle holders are the same size, so it is not obvious where to put thumb or fingers. You might assume that the thumb goes in one of the holes and the other is for the fore-finger or middle-finger if these were to be used for opening and closing like scissors. However, this is not the case. You should use the thumb and ring finger to preserve flexibility about the wrist. In other words, the form (of the needle holders) is not sufficient to indicate their function. Indeed, the function here is not simply a property of the artifact (needle holders) but arises from human capability (experience in using needle holders) interacting with it, and this would make most sense in a specific environment (the needle holders with a threaded needle and tissue to be sutured). Consequently, a concept of affordance must provide "a theory of the world as we experience it in terms of what we can do."[2] We constitute our experienced world through our bodily actions (i.e., we lay down a path in walking),[3] and a theory of affordance should be capable of reflecting this.

So, it makes no sense to speak of an artifact "having" an affordance. And yet, there is an extensive literature, particularly in the fields of design and human-computer interaction in which affordance is treated simply as a property of an artifact. The implication is that affordance is a property that can be designed-into the artifact (say, a coffee cup) and that this property supports, permits, invites, a person to perform an activity (say, drinking coffee). But there is an obvious problem here: the coffee cup also "affords" containing liquid (with additional affordances of transporting the liquid or keeping it hot), cleaning after use, storage, disposal, and so on. Even this trivial example raises the question of *which* affordance we might mean because the same artifact participates in a variety of affording situations.[4] Responding to an artifact depends on appropriate behavioral meanings. A scrub nurse picking up a pair of needle holders, to pass them to a surgeon, might not even use the holes in the handles, but might orient them so that the holes are presented for a suitable grip, or the nurse might handle them differently when picking them up to put them into a tray for sterilizing.

Formal Descriptions of Affordance

One approach that has been explored to capture complementarity (as discussed in chapter 3) in human-environment-artifact systems is through formal descriptions, either in terms of computer programs or statements couched in formal logic. The concept of affordance has received much attention in the field of robotics, and this can be traced to the seminal work of Rodney Brooks and his claim that "world is its own best model."[5] Rather than working from a preprogrammed model of the world in which it is operating, a robot could learn relationships between properties of artifacts and actions that it can perform. From this, "affordances appear from the interaction between the robot and the environment and hence, depend on the world, the agent's motor and perceptual capabilities and its experience."[6] If the available information points to more than one sensorimotor pattern (or if the sensorimotor pattern is associated with more than one set of proprioceptive states) then affordance competition[7] could arise and would need to be resolved. Sridharan and Meadows offer an elegant solution to this problem by separating *affordance relations* from *executability conditions*.[8] What is attractive about these terms is that the focus is on the relations (between elements) rather than solely on the properties of either artifact or agent, and that these relations will change under different situations (or executability conditions). Further, while the work on affordances in robotics makes use of research on human activity, the concepts outlined here could be usefully applied back to discussion of human activity. For example, the emphasis on learned relationships (gained from repeated exposure to similar situations), the need to resolve affordance competition, and the distinction between affordance relations and executability conditions can all apply to human activity.

As we noted in chapter 3 the environment constrains activity in terms of the opportunities it presents, for instance, in terms of artifacts available to the individual or of the interaction between these artifacts and the bodily constraints of the individual. While bodily and environmental constraints can limit or encourage specific actions, the choice of action will further depend on task constraints, which could include a purpose for completing the action as well as some criteria that might define good or acceptable performance.[9] Performance can then be evaluated in terms of elements that constrain or allow activity. This combination of constraints implies a competing set of objectives and a need for multi-objective satisfaction (as we discussed in chapter 2)

Lewin, who provided a precursor definition of what became known as "affordance," developed a simple equation[10] to model behavior (B) as a function (f) of person (P) and environment (E):

$$B = f(P, E).$$

This could be read as a trivial statement that an environment supports an action. But I think the claim here is deeper in that Lewin presents this as a specific person (with defined capability, goals, and so on) responding to a specific environment (with defined features). One can see how this relationship could reflect the notion of bodily and environmental constraints. So, an adult human hand can grasp the handle of a full cup and lift it in a way that a child's smaller hand might not be able to,; in other words, the cup handle (for the adult) can be grasped because its properties (i.e., size and shape) match the disposition of the person, defined by hand-size, and the full cup can be lifted because of the adult's strength (which can also ensure stability during the lift). However, it is worth noting the obvious point here that one can lift a cup without using its handle. In this case, the diameter of the cup defines its property, which is matched by the disposition of the person. Whether one lifts the cup by its handle or not must be influenced by more than the form of the cup. In this case, the cup (and actions performed with it) exist in an environment that, according to Gibson, can be defined in terms of both convention and natural laws. The latter would relate to phenomena such as the temperature of the cup (one might be more likely to use the handle if the cup contained hot liquid) and the weight of the cup (lifting with the handle or grasping the cup above or below the handle would have different repercussions for counteracting the effect of gravity). The former would relate to what might be considered "polite" or "acceptable" behavior in a given environment, such as a greasy spoon café versus the Savoy Hotel.

Formal descriptions reflect relationships between bodily and environmental constraints, but they do not express task constraints or capture the situation in which the relationship arises.[11] Abbate and Bass develop a formal description that reflects situational constraints:

Possesses(affordance$_i$)(Xp, Zq).[12]

This relationship becomes expandable with specific values of the properties, *p*, of an artifact, *X*, which are relevant to a given "goal" and with specific values that define the capability, *q*, of an actor, *Z*, responding to these features. An example they provide involves a cabin door in an airplane. At altitude the door is plugged into its fitting under high external pressure, and on the ground the door can be opened by pulling out a lever and turning it. In this instance, the "environment" is defined in terms of airspace (which could be ground or in flight), aircraft (which could be a specific type), cabin (with properties that change according to airspace, e.g., pressurization), door (which can be opened), lever (which fits into a slot and can be moved). Relations between these can be defined in a formal description as follows:

possesses(doorOpenable)(Xp, Zq) = true if:
Xp. Airspace. Aircraft. Cabin. Door. Lever l [Slot][top_of] = overlapping ^
Xp. Airspace. Aircraft. Cabin. Door.p2 [Cabin][left_of] = contained_within ^
Zq. Airspace. Aircraft. Cabin. Door.q1[position_back] = true ^
Zq. Airspace. Aircraft. Cabin. Door.q1[translate_left] = true.[13]

This description elaborates the context under which the door "affords" opening (in terms of external air pressure and the position of the lever, and in terms of the action performed by the person). What is potentially interesting about this approach is that artifacts may contribute to several "executability conditions," depending on the state of the environment and the goals of the person.

Any formal description is going to be bedeviled by the challenge of completeness; representing *all* contributory features of a situation using formal language quickly becomes overwhelming. Consequently, it is important to ensure that the problem is clearly stated (which, in turn, means that the problem is represented in a way that is amenable to the formal language in which it is being described). While this might offer benefits for verification and validation of design (which is what the formalism is intended to support), it does not provide a plausible model of what humans might do. To be fair, there is no intention on the part of Abbate and Bass that their formal description should reflect human cognition. This is also a way of highlighting the problem that an information-processing view of "affordance" might face: How would the human brain construct the ever-decreasing subtleties of this hierarchical knowledge structure?

To frame this problem more concretely, consider the notion of stimulus-response compatibility (SRC), which has been a staple part of ergonomics for the past half century. To illustrate this idea, in the choice–reaction time paradigm you have a row of four lights in front of you (labeled 1–4), and between you and the lights is a row of four buttons (labeled A–D). The buttons and lights are arranged so that 1 and A are adjacent, and so on. When one of the lights turns on, you must press one of the buttons to turn off this light as quickly as possible. In the adjacent (or congruent) arrangement, when light 1 turns on, you press button A. In an incongruent arrangement, when light 1 turns on, you have to press, say, button C. Not surprisingly, the congruent arrangement leads to much faster performance. Early accounts of the SRC suggested that the performance differences were due to "translation."[14] In information-processing terms, this "translation" is required to allow information-as-content in one "code" (stimulus layout) to relate to another "code" (action).

Contemporary explanations of SRC draw on the ability to extract salient features and pair these with an appropriate response. This is the "dimensional overlap" model,[15] which contrasts the overlap of dimensions (features) in a set (i.e., the congruence of arrangements) with the relevance of features within a set (i.e., how the features of a stimulus relate to a response). The congruent condition has both overlap and relevance. There is much to be said for the empirical evidence from SRC.[16] People prefer arrangements in which the features (light and button) are congruent, and this is termed a "population stereotype" (there is some work to suggest that different cultures might have slightly different population stereotypes).[17] Furthermore, most people produce faster responses with fewer errors in sets of stimulus-response pairings that have this preferred arrangement, as a consequence of SRC. From the perspective of affordance, it could be argued that SRC arises when information from environment (stimulus) relates to ability (response). In other words, there is potential argument that removes the need to appeal to a "translation" or a "dimensional overlap" to explain SRC.[18] Crossing one's hands in SRC experiments leads to an increase in reaction time, even when the position of stimulus and response artifacts remains constant, and this does not seem to be the result of a simple biomechanical constraint; reactions using crossed hands cannot be explained solely by conflict management, as proposed by the dimensional overlap model. This suggests that the relationship between response and stimulus involves more than the predefined mappings that SRC assumes. If we refer to the formalisms outlined earlier, it is difficult to see how these could account for the differences in SRC. In both congruent and incongruent conditions, Xp would be "light on," and Zq would be "press button." So, perhaps, we need to elaborate the Xp description to include Xp1 "light on" + Xp2 "light adjacent to button" (in the congruent condition), and to elaborate Zq1 "associate light label with button label" + Zq2 "press button" in the incongruent condition. But this would produce a near infinite regress in which all possible states need to be defined in order to produce a prediction of an action; in this way, the "action" becomes a matter of matching the features that define a given state. While the descriptions of affordance in this section provide formal descriptions of the relations between human-artifact-environment and their impact on outcomes, I am not proposing that this is meant to describe human behavior. For one thing, formal descriptions like those presented by Abbate and Bass cannot align with the radical embodied cognitive science (RECS) argument used in this book. On the other hand, these descriptions indicate how the elements in the human-artifact-environment system interact, in terms of the functions that are necessary and sufficient to ensure ongoing, reciprocal engagement, and the challenge of defining the "features" that need to be attended to. One reason for including this example is that it presents a set of "objectives" for an affordance. As noted in chapter 2, this set could be collapsed into an optimization problem, where some of the elements are held constant. In the cabin door example, the constant elements (for the human) could be defined by the environment and the physical relations between artifacts, and this would leave a subset of elements that relate to human-artifact interactions. For example, grasp and lift the handle.

Affordance, Capability, and Activity-Relevant Features

If an affordance was a property of an artifact, then, prior to performing an action with that artifact, one would need to "read" the salient features of the artifact that need to be elicited and interpreted. So, the notion of affordance-as-a-property-of-the-artifact could require the extraction of key features, alignment of these features to an appropriate mental model, and use of this mental model to specify the action to perform on the artifact. Having argued against the concept of mental model in chapter 2, we need to evoke a different explanation in order to account for affordances. In subsequent writings, Norman[19] distinguished "perceived affordances" from Gibsonian or "real affordances." Other writers distinguish "simple affordance" (which arises from perception-action coupling) from "complex affordance" (which involves interpretation and response to an artifact's form in terms of the user's culture, history, praxis).[20] Gibson's proposal is that we have a perceptual system that is tuned (through evolution and individual experience) to the environments in which we live. This means that there is no requirement for an interpretive act; we just "see" (or hear or otherwise perceive) a pattern of features to which we can respond: a cup full of steaming hot coffee is "seen" as a different artifact (supporting different actions) than a half-full cup of cold coffee.

As a cat walks through a narrow gap, its whiskers provide sensory information that enable it to keep away from the walls. Humans can judge whether to walk through a gap or whether we need to turn sideways, or avoid the gap entirely.[21] This ability to assess the appearance of artifacts in the world allows us to rapidly judge whether a particular action could be performed in all manner of situations, such as, for instance, step-onto-ability;[22] walking-up-ability;[23] sit-on-ability;[24] step-across-ability;[25] and pass-under-ability.[26] Rapid judgments about whether to turn your body to fit through narrow apertures as you approach these can be made even when bodies have been modified to an unfamiliar size, as when wearing "pregnancy packs," for example,[27] sports shoulder pads.[28] Increasing the weight of the body, such as by wearing a heavy rucksack, can alter judgments of the steepness of a hill.[29] These "body-scaled" perception of features of the environment guide action so that people are able to "see," or become "attuned" to, aspects of the environment in terms of an action that they want, and are able, to perform. What is important in this list of "*x*-ability" (where *x* is any verb) is that it is a reflection not simply of the properties of the artifact, but of the relationship between some features of this artifact and some property of the person. Related to this class of body-scaled affordances are action-scaled affordances; these involve people judging whether an action, such as reachability, is possible, so that people with longer arms estimate artifacts to be close to them[30] or people holding tools estimate artifacts to be closer.[31] These body- or action-scaled perceptions of environmental features provide partial support for the idea of embodied cognition. However, these should not be taken as complete explanations (any more than the suggestion that a cup is "pick-up-able"), because they capture one element (but not all) of complementarity in the human-artifact-environment system.

One might believe that the inexperienced user of an artifact needs to "read" the artifact prior to use to select an appropriate action, while the experienced user simply uses the artifact. However, an artifact such as a pair of needle holders (figure 4.1) might not obviously yield to such a reading without an understanding of the action involved in its use. Trial and error *might* allow you to try different grips (perhaps you could perform some basic suturing activity using the wrong grip without realizing that there is a superior means), but feedback would be required to confirm that the action was appropriate. This raises the question of what information is required to know how to *simply use* artifacts?

For this book, the challenge is to provide an account of what it means to "simply use" something. In connection with the contrast between "reading" and "simply using" an artifact, Humphreys[32] offered two distinct routes from artifact to action: (1) perception of specific features of an artifact (which Humphreys calls a "structural description"), which can be associated with knowledge of how to use that artifact (which he calls an "action description"); and (2) a direct link between the structural description of an artifact and the action description of how to handle that artifact (which he calls "affordance"). Notice that route 1 echoes Norman's "perceived affordance," and this is distinguished from route 2 affordance. This distinction between route 1 and route 2 hints at the question of what information is used in “reading” an artifact and how this information is obtained from the environment. When Norman uses the word “perceive,” it is not in the same manner that Gibson uses it. Norman regards perception as an active process of assigning meaning to an artifact’s features and associating that meaning with an appropriate action. For Gibson, perception is sensitivity to information that corresponds to action; there is no intermediary process of interpretation. This further illustrates the contrast between “information-as-content” and “information-as-context” and highlights a fundamental aspect of affordance.

Assuming that an artifact is perceived in terms of information-as-content, it might offer competing interpretations, and the challenge is how this competition might be resolved. We could, following Norman’s proposal, seek to “read” the artifact so that certain of its features align with the function that the artifact can support. In this way, the size of the cup and the position of its handle might be more salient than the pattern that is painted on it. The idea of “reading” an artifact might be appropriate in disciplines such as archaeology, particularly when the form of the artifact is so unfamiliar that it does not easily support “reading,” but it does not feel as if this is something that we perform with familiar artifacts in our everyday life. Indeed, according to his material engagement theory, Malafouris argues that understanding the manner in which people physically interact with the artifacts[33] is necessary to understand the embodied nature of interaction between person and artifact and how this supports ongoing, reciprocal engagement. If this is the case, then the range of affordances needs to be considered not solely in terms of the artifact but also in terms of the situation in which the artifact is placed, where the situation is shaped by the environment and the capabilities of the human. This idea has been developed by Rietveld and his colleagues in the skilled intentionality framework in their proposal that there are “landscapes of affordances.”

Skilled Intentionality Framework

The skilled intentionality framework (SIF)[34] has four basic premises, which align neatly with RECS:

1. There is no division between “higher” and “lower” cognition; both can be understood in terms of skilled activities of engaging with situations in the world.

2. Skilled activities are temporally extended processes in which agents coordinate to multiple relevant affordances simultaneously.
3. The affordances the environment offers are relative to the abilities available in a form of life.
4. “Higher” order cognition does not necessarily depend on mental representation.

SIF accounts for an “individual’s selective openness and responsiveness to a rich landscape of affordances.”[35] Broadly, SIF seeks to explain how people can encounter multiple potential affordances in the environment and selectively respond to those that are salient to a specific situation.[36] In SIF, intentionality is considered in terms of skillful coping. In this manner, “Affordances are relations between aspects of a material environment and abilities available in a form of life,”[37] which can offer opportunities for action. In order to explain how we selectively respond to affordances, SIF proposes that there are “multiple simultaneous states of action readiness for engagement with affordances.”[38] An interesting position that the developers of SIF take is that “skill” is not restricted solely to physical activity but can encompass all aspects of behavior, including those that we might term “cognitive” or “cultural.” This accords with Ingold’s notion of skill as “*the capabilities of action and perception of the whole organic being (indissolubly mind and body) situated in a richly structured environment*.”[39] This notion is important because it allows us to move beyond so much of the discussion of “affordances,” which have a tendency to focus solely on physical activity.

SIF proposes that affordances can be considered in terms of “solicitations,” in which those affordances that are relevant to a given situation are preferentially attended to. This borrows from Gestalt thinking (in its implied meaning of “invitation character”) and Merleau-Ponty’s concept of “intentional object” to suggest that the “invitation character” of the artifact aligns with the lived experience of the skilled actor. These ideas were familiar to Gibson, as he had lectured on Merleau-Ponty’s work.[40] Solicitations allow the human to have “maximal grip” (in Merleau-Ponty’s terms) in their engagement in the situation. From this perspective, affordances are relations between aspects of the ever-changing sociomaterial environment and the abilities available in a form of life.[41] In this case, a “form of life” is an expression coined by Wittgenstein[42] as a means of describing the routine or patterns of activity of our workaday and everyday lives. Thus, artifacts are not simply physical things but also value-rich ecological objects (where the "values" are defined by the social and cultural milieu in which they are encountered, i.e., by Gibson's notion of "conventions"). In this way, skillful coping is not simply the enactment of physical activity but also social and cognitive behaviors that are possible and plausible for members of a given community. Situating affordance within the social setting means that activity can be socially constrained in terms of what is an acceptable way to employ an artifact is, as noted in chapter 3.

The Politics of Affordance

The linkage between environment and action is influenced (according to Gibson[43]) in three ways:

1. convention;
2. projection (arising from the effects of physics, such as a shadow behind an artifact), and
3. natural laws.

"Convention" and "natural laws" provide the twin poles from which an epistemology of artifacts can be defined. Convention represents sociocultural norms, while physics reflects the world as it is (Gibson assumes an objective reality that can be reliably defined through the laws of physics). From this, Gibson sought a "lawful" relation that would define "affordance." For Gibson, affordance "*implies the complementarity of the animal and the environment*."[44] "Complementarity" can occur only in the interactions between "animal and environment." In other words, "affordances" exist in the relations between features of the environment and the capabilities of the animals in those environments.[45] In chapter 3, we discussed how an environment can be considered in terms of the ecological niche for a type of animal. The concept of affordance develops this further, in that it is concerned with the ways in which the ecological niche can be considered in terms of its activity-relevant features and how these can be responded to by the animal. In the words of Gibson, "*A niche is a set of affordances*."[46] That is, sensory capabilities of the animal become linked to specific features of the environment that constrain (or support) that animal's actions. When the animal experiences a similar situation, then that action will be more likely to be performed. As Pickering says, "*Affordances . . . are the behavioural meanings of the environment for particular organisms*."[47]

While the "behavioural meanings" of affordances change with the "webs of relations" between the artifacts, humans, and the environment, much of the previous research on affordance has focused on relations between artifacts and their users. But, as the SIF highlights, this omits the importance of the ecological niche in which the landscape of affordances exist, and, in particular, the social dimensions in which normative action is defined. Gibson spoke, for instance, about "convention" as one of the aspects that influence complementarity, but the literature has been surprisingly quiet on this. As a consequence, the manner in which affordance relations might change in different social settings has not received as much attention as it ought to. The recent book by Jenny Davis[48] not only highlights this omission but also provides substantial contributions as to how to conceptualize and address these problems.

Davis begins with the observation that the literature on affordance often assumes a binary distinction between having or not having affordance. We have noted how the SIF has shifted debate beyond this to a more nuanced sense in which "skill" (as adaptive coping) provides a way of conceptualizing the ability of an individual, both in terms of the ways in which actions are performed and also in terms of the ways in which goals or intentions are defined. However, the manner in which such goals are made meaningful to individual actors is less clearly developed, and this where Davis makes a key contribution. For Davis, affordances involve mechanisms through which and conditions under which they are effected.

In terms of mechanisms, Davis draws a loose distinction between "bids by" and "bids on" an artifact (although, of course, the term "artifact" ought to be read in terms of the relations between person-artifact-environment). In the category of "bids by" she includes "requests" (which are, I think, synonymous with "solicitations" in SIF and Gestalt-inspired versions of affordance). However, she elaborates on this category with the inclusion of "demands." For example, if you want to prevent people from walking into a particular space, you could string a piece of rope between poles as a "request" to them to avoid this area, or you could build a metal fence as a "demand" for them to keep out. She makes the interesting point that it is not simply the physical property of the artifact that implies a demand; for example, the police tape that might be placed around a crime scene has the flimsiness of the rope but a much more forceful message in terms of its meaning. Thus, the difference between request and demand is a matter of social convention as much as physical form. So, a fence could be viewed as a request to protestors who desire to occupy the area beyond it.

In the category "bids on," she includes affordances that "encourage" behaviors—for example, large plates enable people to have large portions of food—or that "discourage" behaviors—for example, the character limits on Twitter discourage long-form content. This is not to say that, in either case, behavior is prevented (you could put less food on a large plate, or you could post multiple tweets to form a long message), but there are, perhaps, additional social constraints or physical demands on countering these bids on the artifact. At the extreme ends of these constraints or demands are affordances that "refuse" an action. She uses the example of Robert Moses's civic planning in New York in which low bridges that could not allow public transport to pass under them were placed across the rivers into some boroughs. Or affordances could simply "allow" an action without commitment, as in the case of a fork in the road.

What is apparent in her expansion of the types of affordance is that Davis emphasizes the social, moral, ethical, political, and other values that inform particular stances that are taken in the design of artifacts. Her view of the conditions of affordance can be expressed as "How does this object afford (mechanism) and for whom and under what circumstances (conditions)?" These conditions are covered by perception, dexterity, and cultural and institutional legitimacy. In this context, "perception" is akin to Norman's "perceived affordance," in which the artifact's "meaning" depends on the awareness of its users and their interpretation of features. To continue with an embodied cognition (rather than information-processing) argument, I would suggest that "perceived affordance" could be reflected by Merleau-Ponty's notion of an "intentional object" (discussed in chapter 3), which includes that collection of features in an artifact that correspond to a given intention for its user. Implicit in this notion is that different users, by virtue of their experiences of the world, experience artifacts as different "intentional objects." For Merleau-Ponty, we are embodied perceivers who act upon the world (and have the world act upon us). Sensations are the basic material of perception, not as a stream of inputs that need to be processed, but as a pattern against a background. We perceive those sensations to which we are most tuned. "*I discover vision, not as 'thinking about seeing,' to use Descartes expression, but as a gaze that grips with a visible world.*"[49] Thus, each person seeks to exert *maximal grip* on the world through responding to the pattern that provides an optimal collection of sensations for *that* individual, with *those* capabilities and goals, in *that* situation. As one develops experience of things in the world, so one begins to respond less to the specific features of each individual thing and more to essential aspects of the situations in which these things exist. Key to Merleau-Ponty's phenomenology is the notion of intentionality, which is concerned with how we "see" an artifact in terms of how we intend to interact with it (rather than as a collection of features). That is, we see the intentional object in relation to our goal. One way of appreciating this is through the concept of "Gestalt" (with which Merleau-Ponty was familiar), which is not some property of the artifact but rather the combination of the sensory stimulation evoked by an artifact in a given context. This means not only that the Gestalt is more than the sum of its parts, but also that the artifact can be interacted with differently under different conditions.

Merleau-Ponty's notion of "essential aspect" might, at first glance, feel overly metaphysical and clearly at odds with Gibson's arguments about the physical properties of artifacts. However, another reading of "essential aspect" would be to consider the specific set of features that the person sees (e.g., when the artifact is viewed from a particular angle) as defining the information that constrains an action. In this case, the visual appearance of the artifact will be influenced by "physics" and the action by "convention" (to return to Gibson's explanation). If one takes this point a little farther and rephrases this as the specific set of features to which a person attends, then it is possible to see the artifact not as a single, homogenous entity but as collection of features that can be attended to for different actions. So, the situation in which we look at an empty cup has a different essential aspect to one in which the cup is full of steaming coffee. In both cases, the artifact provides opportunities for action (which is another formulation that Gibson used to describe affordance), but the actions depend on the set of features (and on the interactions between these features and the person and the situation in which the action is to be performed).

A second condition of affordance for Davis is dexterity. This relates to the ways in which degrees of freedom (chapter 1) are managed and to the skillful coping of the individual. A key issue for the politics of affordance relates to the question of what dexterity involves. For people with a visual, physical, or other disability, skillful coping and dexterity involves their ability to adapt to the demands of the environment. Affordance, in this context, should mean the adaptation of the environment and artifacts to better support and enable their dexterity. Bad designs can exclude users. For example, the designer could exclude swathes of users on the assumption that everyone has the same abilities. While design in practice focuses on the needs of specific user groups, there can still be challenges for people at the extremes of these groups. More insidiously, Davis points out how designs can exclude and marginalize potential users in ways that are due to more than just anthropometry or physiology.

Davis's third condition of affordance involves cultural and institutional legitimacy. Earlier in this chapter, I made the trivial observation that picking up a cup to drink by grasping the handle or the rim might depend on the environment. More precisely, Davis argues that "as a condition of affordance, cultural and institutional legitimacy addresses the way one's location within the larger social structure and the related norms, values, rules and laws of a social system inform human-technology relations."[50] While Gibson's notion of "convention" nodded toward this condition, his work (and much of the subsequent debate surrounding affordance) has not grappled with the implications of the ways in which such conventions reflect the power balances within society—by economics and access to technology, for instance, as much as norms of "good" behavior. To a great extent this is due to the fact that the "environment" is not simply physical but also defined by social conventions and normative practices (as recognized by SIF). This means that the normative behaviors that an artifact is intended to support need not apply in all situations. Consequently, the concept of affordance has to cover the artifact, the ability of the user, the normative social conventions, and the physical environment in which actions involving the artifact are performed. In the next section, I consider how affordance relates to information and how this influences interpretation of artifacts.

Affordances as Information

For Gibson, "information" is available in the relations between features in an environment and this results in an action, but this information can arise only for an agent attuned to it. In chapter 1, I introduced the distinction between information-as-content (which can be processed and assigned meaning) and information-as-context (which influences action). In this section, I relate this distinction to the concept of "affordance." Koffka wrote that artifacts "tell us what to do with them"[51] through their "demand character."

Interestingly, Koffka was a colleague of Gibson's in the 1930s,[52] and Gibson quotes his words as a description of "vivid and essential features of the experience itself."[53] Koffka claimed that artifacts solicit actions: "A fruit says 'Eat me'; water says 'Drink me'; thunder says 'Fear me.'"[54] From the tradition of American pragmatism, Peirce[55] suggested that the artifact (as a specific thing) can be perceived as a representamen (as a class of thing), which is then made sense of as an interpretant (as a concept). So, when Koffka wrote that an artifact "says" to do something, he is treating the artifact as a representamen in order to evoke an interpretant. But even a cursory consideration of an artifact, say a specific item of fruit, could take different forms, such as an apple on a tree, in a fruit bowl, in a lunch box. Each of these can be perceived as different a representamen (e.g., nature, still life, food) and, in turn, can result in different interpretants (e.g., harvest, painting, eating).

To claim that what something (or someone) "says" has a single meaning is to collapse a host of potential representamens into a single interpretant. This would imply that (a) there is only one meaning possible, (b) that one has little choice but to perform an action in response to this meaning, and (c) that the artifact requires a specific action (because the word "say" has an imperative force that arises from the artifact itself). For some approaches to design thinking (as discussed in chapter 2), this might help to bound the solution space for the problem being solved by a focus on the form of the artifact. But I doubt if any designers would agree that this all they do.

We could separate affordance from the interpretive act. Gaver[56] introduced terms such as "false affordance" (in which the form of an artifact implies a possible, but undesired, action; say, a decal on a product that looks like a button you can push) or "hidden affordance" (in which information is obscured and needs to be discovered). The notions of false and hidden affordance imply the need to search for and make judgment on the information that can be perceived in an artifact. This takes us further into the realm of information-as-content. Both Norman and Gaver present a theory in which the function of the artifact needs to be read from the form of that artifact. While it might be useful to consider the consequences of "false" or "hidden" information, this confuses the definition of affordance. We could take their ideas to mean that we should be concerned with designing visual signifiers that cue an action; but surely this is much the same as stating that the form of an artifact signifies its functions? If affordance is to provide a different and useful perspective on design and use of artifacts, then it needs to be more clearly articulated. To return to the example of scissors versus needle holders—one could place the thumb in the smaller of the two holes and then see where the fingers were positioned. In this respect, the "meaning" of an artifact could just as easily be physical rather than cognitive. This implies that there could be a "semiotics for action" in which, rather than reading the meaning of an artifact, we respond to its potential for action. In this case, the features of the artifact to which we respond could be influenced by task and environment (as an "intentional object") and the "meaning" of these features would be defined by the salience of the artifact to our skillful coping.

There is an implication that salience is agent-specific, and so, information itself could be agent-specific. Worried that defining such relations in an agent-specific manner would mean that there could not be natural laws, Gibson wanted to define information in generic terms. In the 1930s, the Gestalt psychologists developed concepts that directly inspired the notion of affordance. Kurt Lewin introduced the term *aufforderungscharaktere*[57] (which can be translated as "invitation-character" or "prompt-character," or "demand character" as per Koffka) as a way of indicating how the character, or property, of an artifact solicits a certain action by a person. For Lewin, behavior was influenced by "valences," which were modified by the person's motivation, and situational demands. For example, a cup of coffee might *not* afford drinking if the person was trying to reduce caffeine intake or was in a rush to leave for a meeting or was angry or was concerned about the cleanliness of the cup in which the coffee was served. Thus, the *aufforderungscharaktere* is not simply an attribute of the artifact but the sum of factors that define the situation in which the artifact is encountered.

The idea of a collection of elements that corresponded to valences was antithetical to Gibson's views. For Gestalt psychologists the artifact has a phenomenal value that the person experiences. Gibson wanted to claim that these phenomenal values are redundant. Rather, one ought to focus on the physical property of an artifact as something that can be directly perceived and, through this direct perception, acted upon:

> The affordance of something does *not change* as the need of the observer changes. The observer may or may not perceive or attend to the affordance, according to his needs, but the affordance, being invariant, is always there to be perceived. An affordance is not bestowed upon an artefact by a need of an observer and his act of perceiving it. The artefact offers what it does because it is what it is.[58]

Here, Gibson argues that the affordance must be an invariant property of the artifact, and this property must be independent of the observer. This would constitute the "natural laws" that underpin affordance. In terms of natural objects, likes pebbles or sticks, one *might* wish to claim that the physical appearance of the object would remain constant irrespective of where it was found or who found it. If this is the case, then Gibson seems to be arguing that one *can* design an artifact to have a particular affordance.

And yet, Gibson continually contradicted this assumption of invariance. This is especially the case when considering manufactured artifacts. In this quotation, about a post-box, Gibson is arguing both for and against the experienced reality of the artifact:

> To be sure, we define what it is in terms of ecological physics instead of physical physics, and it therefore possesses meaning and value to begin with. But this meaning is meaning and value of a new sort. For Koffka it was the phenomenal post-box that invited letter-mailing, not the physical post-box. But this duality is pernicious. I prefer to say that the real post-box (the only one) affords letter-mailing to a letter-writing human in a community with a postal system.[59]

First, the question of how one might separate a "phenomenal" from a "real" post-box is tricky. Possibly one might consider the "real" post-box as a physical artifact in the street but it remains a phenomenal post-box by virtue of the fact that its "salience" arises in relation to a specific activity (letter mailing) arising from a specific human capability (letter writing) within specific sociocultural norms (community with a postal system). Of course, the post-box could "afford" other activity—one could drop trash into it (which would contravene norms), for instance, or one could use it as a table, perhaps to read a map or check the contents of one's purse, or one could use it to hide behind. By wanting to argue against the Gestalt notion of valence, Gibson seems to be forcing the concept of affordance onto the artifact itself rather than from the web of relations in a given situation (which contradicts what he has said elsewhere).

Can Affordances Be Designed?

If an affordance is *not* a property of an artifact, how can affordances be designed? In order to address this, I want to introduce a related concept, that of "habitability," which refers to "the match between the language people employ when using a computer system and the language that the system can accept."[60] The analogy I wish to draw between habitability and affordance relies on the assumption that the action that a person performs (whether physically using an artifact or speaking to a computer) involves the person's best effort after salience in a given situation. For habitability, action is constrained by semantics (the goal that the user seeks to achieve), dialogue (the history of utterances to that point), syntax (the structure of spoken commands that would be acceptable), lexicon (plausible words that relate to the user's goal), and recognition (performance of the speech recognizer). In speech recognition systems, the user could be given explicit instructions in terms of words to say or could be given an open-ended prompt, as in "say in your own words what you would like help with." From our constraint model of habitability, we suggested that identifying the most likely constraint that contributed to an error could help in providing useful guidance for the next action. For example, if the user selected a word that was not in the lexicon, the computer could suggest a word from its lexicon, assuming that it was able to make a guess at the user's goal; if it could not guess that user's goal, it could offer a set of possible "goals" for the user. From this, the computer provides guidance on the basis of gaps between its expectation and observation of user action. In a similar manner, "seamful design"[61] seeks to identify the gaps or "seams" between different functions of a computer as a resource to guide user action.

Rather than considering how the form of the artifact will solicit an action, it is more important for the designer to think about the situation in which different users could respond to the artifact as an intentional object in terms of their dexterity and skillful coping. Doing so involves consideration of the affording situations in which the artifact might participate—that of the human-artifact-environment system and how these could generate different activities and outcomes. Several of the examples of design practice that were explored in chapter 2 share an affinity with the open-endedness of this approach. Sketching and model-making become the basis not simply for exploring what the artifact might look like but, more importantly, for considering how people will respond to it, how they might interact with it, and what could be done to encourage activity in particular situations. In this respect, affordances are contingent on environment, situations, and goals and dexterity of the users, as much as the form of the artifacts. In other words, affordances arise from the dynamics within the human-artifact-environment system. In the next chapter, I elaborate on this point and explore an approach to design that is closely allied to Gibson's ecological psychology and, I argue, that allows the principles of RECS to apply to design.

5 Ecological Interface Design

Introduction

The practice of ecological interface design (EID) originated in the work of Jens Rasmussen, who spent much of his career as a cognitive systems engineer for the Halden Nuclear Reactor Project in Denmark. Over the course of his work, he studied a range of issues, from the ways in which people diagnose faults in electrical circuits to how they monitor and control complex industrial processes (figure 5.1). Across this work he developed a theory of human activity (focusing primarily on industrial processes but extensible to any situation in which people interact with technology). As we shall see, this theory owes much to the concepts of Gibson; even if Rasmussen's theory does not directly draw on embodied cognition, there is much overlap and scope to build on the parallels.

Given the prevalence of digital technology, the design of the user interface (i.e., the manner in which information is presented to the person and the manner in which the person's activity is invited and supported by a device) is of central importance. In human-computer interaction (HCI), an influential account of the reciprocal relationship between technology and human activity is the task-artifact cycle[1] (figure 5.2).

The task-artifact cycle points to two core themes for HCI and for design more generally. The first is that the tasks that people perform are constrained by the device. The usefulness of this apparent truism becomes obvious when you realize that "task" is not simply the performance of an action but also the setting of a goal and definition of fitness (i.e., how well the goal is met). In this respect, the task-artifact cycle illustrates the ongoing reciprocal engagement in the human-artifact-environment system. It does so in terms

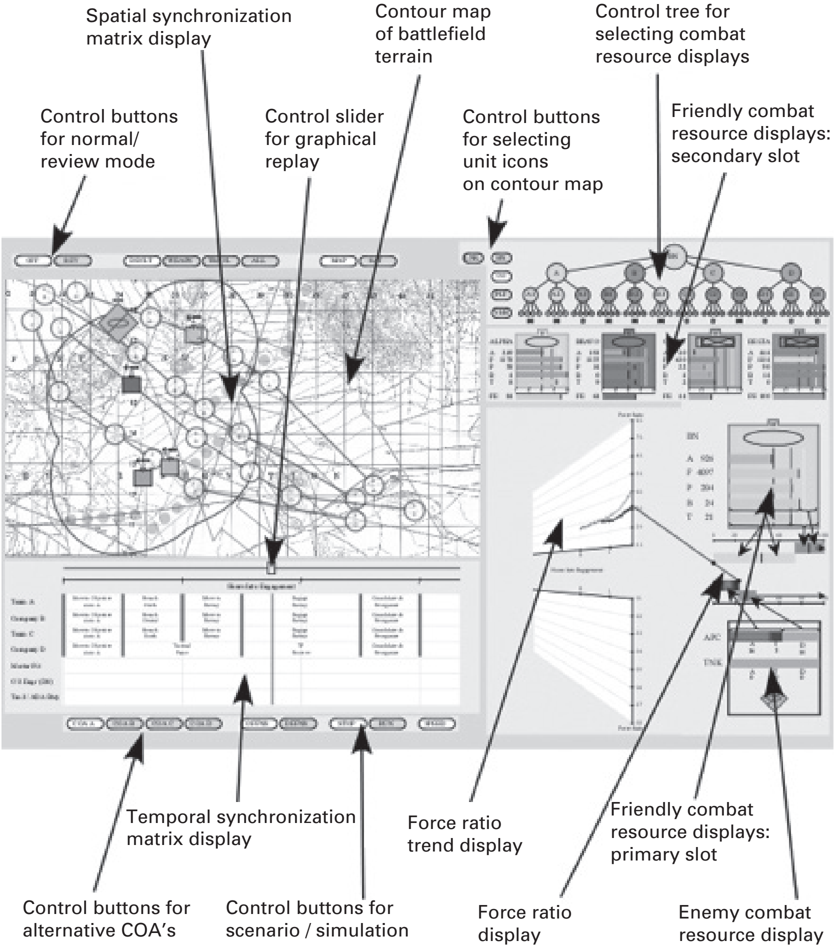

Figure 5.1
Example of an ecological interface.

of the relationship between human and artifact defined as action (task) and outcome (goal). While the environment is not explicitly included, one can expect the outcome to be further constrained by the environment. Further, the focus on "task" might imply solely physical action but it carries a richer meaning that includes the interpretation of the goal. To illustrate this, consider writing a letter with a typewriter (if you can remember such a device), or a word processing program on a laptop, or with a stylus on a tablet. There are obvious differences between these artifacts and in the physical

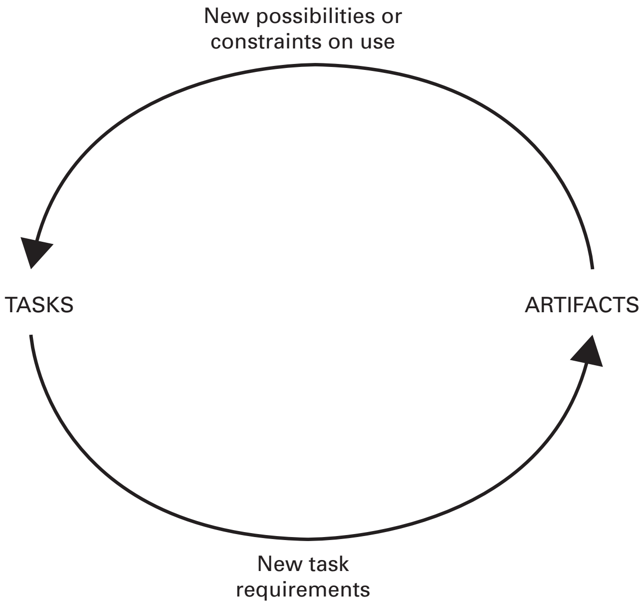

Figure 5.2
Task-artifact cycle.

actions that are made to form words and to manage the layout of the pages of text. In addition, there are differences in how mistakes might be corrected, which could have an impact on how much care you might apply in forming the words (particularly if, like the typewriter, the technology is unforgiving of error). In this case, the performance of each task could be made with greater or lesser awareness of contingent tasks (such as error correction). Moreover, the "outcome" of the activity (in terms of the content, layout, and appearance of the letter when it is printed) has a bearing on the performance of the tasks, both in terms of the meta-tasks associated with tabs, margins, spacing, and other settings for the page setup, the choice of font, and son on and the impact of these meta-tasks on the activity (e.g., should these meta-tasks be performed prior to typing the letter or can they be amended at any point during the typing?). These considerations point to the second theme, which is that (re)design of the artifacts can be based on problems that arise from these contingent and meta-tasks. As these "problems" become more clearly defined, so their "solution" leads to new versions of the artifact.

Depending on how one views this relationship between problems and solutions (in the development of technology), one can conceptualize the design of new artifacts either as a linear progression towards the "best" version or as a patchwork of changes. Although there is a sense in which the problem-solution pattern might imply a dialectical progression, few people subscribe to the idea that technology develops in a linear manner. Rather, most accounts of the evolution of technology recognize that developments are in response to specific problems (and different people using the same artifact for different tasks might encounter different problems). In an interesting account of technology development, Bijker[2] suggests that the recognition of "problems" with technology relate not only to the people using it (and their uses of it) but also to the "technological frame" that they apply to it. For designers, the "technological frame" relates to the tradition in which you have been working and to the manner in which this tradition can be applied, through analogy, to new problems. For example, in the development of the Internet, some researchers came from a telephone network tradition and emphasized the ways in which data could be combined into packets to route through the network; others came from the Air Defence Systems and emphasized the need for the network to be secure and robust; others came from traditions that emphasized human interactions with information.[3] Each tradition brought its own definition of concepts such as "network" or "information" and different expectations of how people would interact with these. One of the key aspects of the development of the Internet (in addition to the technical achievements) was agreement on which aspects of which definitions to include in order to reach consensus on what to build and how it should operate.

The task-artifact cycle illustrates the manner in which user activity is constrained by a given technology. This concept of constraint is not only critical to the definition of embodied cognition presented in this book, but also key to the approach of EID. In order to explain this approach, we need to appreciate the relationship between the idea of task ecology and cognitive work analysis.

Cognitive Work Analysis

The notion of task ecology was presented in chapter 3. Rasmussen proposed several ways to support a description of the ecology in which a system operates. For example, in AcciMaps[4] the system is decomposed into levels that range from the political/regulatory to specific activities performed by people. The aim of this description is to provide a framework within which to understand the interacting constraints on a system during its operation prior to an accident (hence, the "acci" in the title).

From the notion of task ecology developed in chapter 3, we can say that constraints, on the human-artifact-environment system can be defined at different levels, such that there is the level at which certain activities are impossible, there is the level at which certain activities prevent performance of other activities, there is the level at which certain activities could result in dangerous or undesirable outcomes, and there is the level at which some outcomes are valued more highly than others. In other words, constraints will be defined at different levels, from physical to social and economic, to political or societal. Rasmussen describes these in terms of whether responsibility for outcome lies with the operators, their manager, or the government agencies regulating the industry. In this broadening of the definition of a system, Rasmussen was considering the social implications of technology, with much the same motivation as socio-technical systems[5] or actor-network theory.[6] In Rasmussen's approach, each level has a bearing on the salience of information (i.e., sets of features from the environment) or the choice of action or the definition of outcomes (in terms of acceptability or desirability). The different levels give rise to the idea of a part-whole configuration, which reflects the focus of attention (from physical to social) of the people acting in that environment. An obvious consequence of this idea, and one that features in accident investigations, is that undue focus on one level could result in failure to recognize, or misinterpret, constraints at another level. While I have used the word "levels" in my discussion of Rasmussen, it is important to note that he was more concerned with "transitions and overlaps" within a hierarchy of "units" (or views) of a work domain. Other methods developed by Rasmussen and his colleagues focused on the boundaries of the system and are couched in terms of what the system seeks to achieve and avoid—in other words, on the relationship between the functional purpose and desirable outcomes. This is presented in the form of a work domain analysis.

In the work domain analysis (figure 5.3) means-ends relations (as considered in terms of problem-solving in chapter 2) are captured in a hierarchy of five levels (from bottom to top): (1) the physical form, which

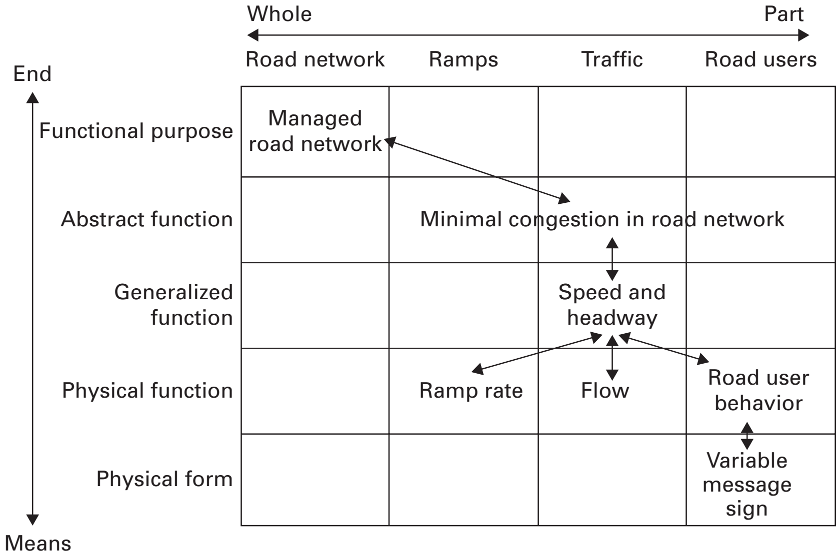

Figure 5.3
Work domain analysis.[7]

can be considered in terms of the features of the environment (in terms of the input to the system, e.g., control devices, information sources, and the like); (2) the physical functions that can performed by the system on the environment; (3) the generalized function, which can be considered in terms of the high-level goals which can satisfy the abstract functions; (4) the abstract function(s), which can be considered in terms of the desired outcomes ("values and priorities") that constrain system activity; and (5) the functional purpose which is the raison d'etre of the system.

It is not my intention here to provide guidance on how to create this diagram (but there is excellent guidance available for the methods that underpin this general approach of cognitive work analysis [CWA][8]). The "abstraction" reflects Rasmussen's observation that people (particularly those he studied who were fault-finding in electronic circuits or controlling power stations) engage in different forms of cognitive activity at different levels, from the immediate local features of the environment to more global concerns as to the functional purpose of the system. The "decomposition" reflects Rasmussen's observation that the different "abstraction" levels will define salience in different ways. This moving between levels of detail is also characteristic of domains such as intelligence analysis, where analysts appear to move between "broad" and "narrow" views of a situation.[9] Taken together, the notions of abstraction and decomposition suggest that problem-solving involves consideration of a system at different levels of purpose and in different levels of detail.

The abstraction-decomposition space can be "read" top-down (in terms of how the desired outcomes enable and constrain activities) and bottom-up (in terms of how the objects in the system support the performance of activities). From figure 5.3, the description is concerned with appreciating the "why" of the system (in terms of what it is seeking to achieve and how success is defined) and the "how" of the system (in terms of what activities contribute to success). In this way, the abstraction-decomposition space illustrates the relations between the physical function and functional purpose levels, and it *appears* to reflect the knowledge that is used in problem-solving. Having said this, I would add that it is not apparent that the abstraction-decomposition space or the analysis that produces it is capable of capturing "tacit" knowledge, so much as assuming that the analyst has produced a "formal" description, which returns us to the discussion of information-as-content versus information-as-context in the earlier chapters. This point will be explored further in the next section. For this section, a further point to note is that the abstraction-decomposition space implies that means-ends analysis relates to the connections between each level. This means that the definition of a "means" slips from an activity to a "function" in ways that are not obviously useful or ontologically reliable (similar problems relate to the conflation of "end" with "purpose"). This conflation of "means" with "function" can create problems, particularly for the novice user of the method and, as far as I am aware, is not an issue that the CWA community has fully resolved.

Before considering decision ladders, it is worth digressing slightly on the topic of "tasks." Ergonomics has long been concerned with defining the "tasks" that people perform in the workplace. In these terms, a "task" is a discrete action that is performed in order to achieve a goal. From the perspective of CWA, however, this notion is too restrictive because it implies that the task-goal mapping points to a "one-best-way" to do work (i.e., the criticism points to a Taylorist tendency to routinize work to the most efficient strategies). I do not agree with this interpretation (task analysis methods are quite capable of reflecting alternative strategies), and a "task" is essentially a system goal that needs to be achieved (this echoes the point made earlier regarding “tasks” in the task-artifact cycle).

Decision Ladders and Decision Strategies

We began this chapter with a consideration of “tasks” (in the task-artifact cycle) and the observation that, by implication, the “task” includes its goal and definition of fitness. Navigating the abstraction-decomposition space involves focusing on a task as a possible route within the space of possibilities—in other words, choosing a course of activity (defined as a “strategy” by which information is selected and actions are performed) that relates to a functional purpose. A strength of CWA (in comparison with most other methods) is that it can comfortably accommodate multiple alternative routes. This means that, rather than being bound to a specific instance of the system in a specific instance of an environment, the approach seeks to be event-independent. To do this, CWA uses a form of visualization known as the “decision ladder.”

Schematically, the decision ladder (figure 5.4) represents activity as a simple, linear input-output flow. The input is a stimulus from information sources (described as physical objects in the abstraction hierarchy), and the output is an action. The decision ladder assumes that there is a canonical sequence of steps, which can be considered in terms of information-processing psychology. Diagrammatically, the flow is split into two “legs”; in the left-hand leg, the input leads up to the goal (purpose) as a process that identifies a discrepancy between input and desired state, and, in the right-hand leg, an action is specified that returns the system to its goal state. Each of the intervening steps is labeled in terms of the information sources that are required to complete the task. In essence, the process is one in which the environment is sampled and a step completed (either as an action to collect information or to process the information to pass to the next step). From figure 5.4, one clearly sees the means-ends approach to problem-solving and the information-processing tradition. The reader might be a little nonplussed as to why I am presenting this. The step ladder clearly shows the serial, “production line” of information-processing that I complained about in chapter 1. Also, it seems to reflect the standard operating procedure (SOP) by which activity is performed. As is well known, experience leads us to adopt all manner of paths to circumvent the SOP,

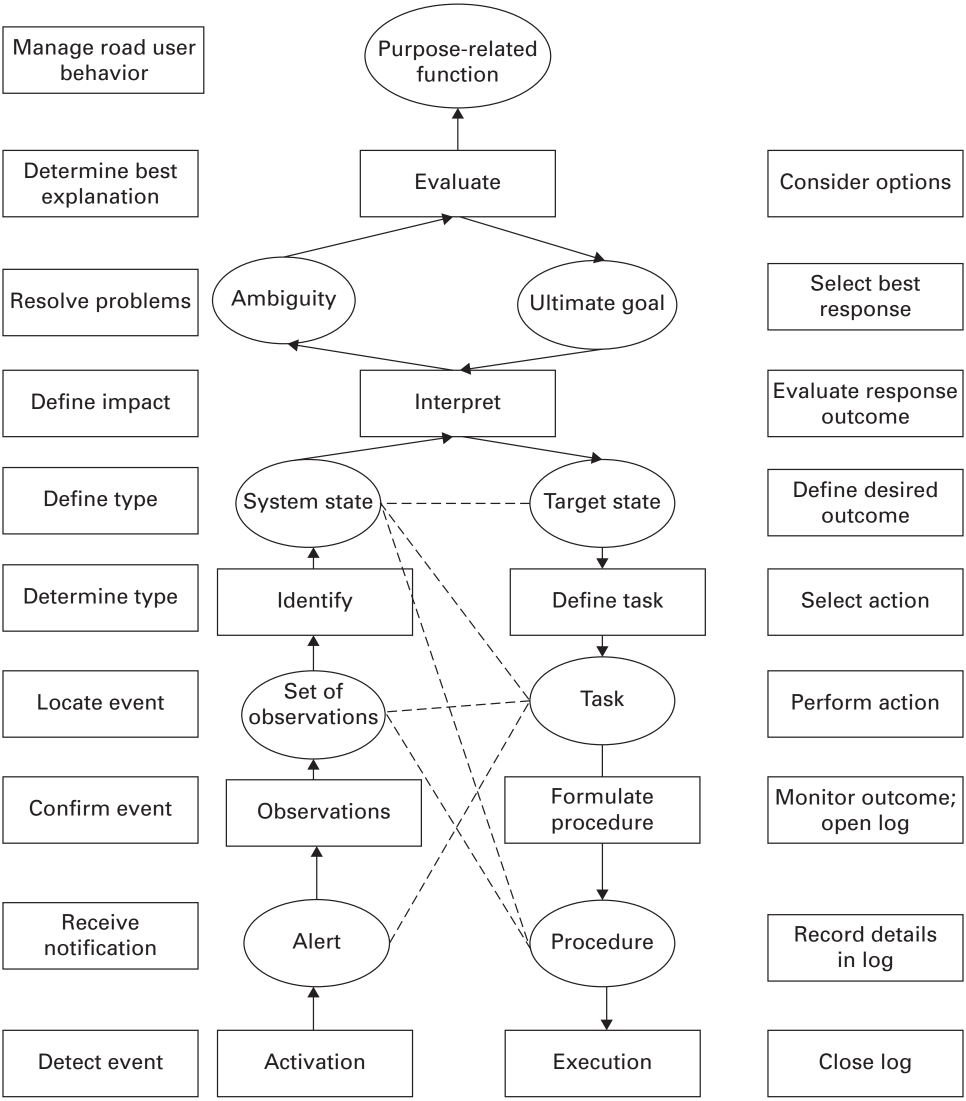

Figure 5.4
Decision ladder.

and it is these paths (shown as dotted lines on figure 5.4) that are intended to capture the strategies that are used to respond to different environmental situations or system states.

Lintern[10] proposes that the paths between each side of the step ladder can take three forms: state transitions (a term he prefers rather than "short-cuts"), which retain commitment to the SOP in that these might be performed under certain states; shunts, which involve deliberative transformation of one "state" to another; and leaps, in which there is a direct association between "states" with no intervening process. While I am using the word "states" here to refer to the human-artifact-environment system, it is not clear how Rasmussen uses the word. For Rasmussen, the decision ladder shows how "rational, causal reasoning connects the 'states of knowledge' in the basic sequence."[11] It is a moot point as to whether these "states of knowledge" (including the sequence itself) constitute a mental model (internal representation) that the person creates and maintains during their activity. An information-processing interpretation would assume this to be the case. Alternatively, Rasmussen saw the decision ladder as a "a map useful to represent the structure of such a model"[12] (which could, I guess, be read as the SOP by which a trained person might conduct their work). Yet, as we shall see later in this chapter, his use of ecological psychology could mean that the "states of knowledge" are just as likely to be realized with objects in the environment (or, at least, objects on a visual display). This seems to echo the notion of off-loading from distributed cognition and the separation of knowledge in the world and knowledge in the head.[13]

In terms of the strategies that can be used to navigate the space of possibilities, Rasmussen contrasted "skill-based," "rule-based," and "knowledge-based" action. In "Knowledge-based" action the problem solver seeks to define the problem and deduce the most appropriate sequence of actions by which to solve the problem. This follows from the information-processing approaches to problem-solving that we saw in chapter 2 and is intended to reflect the strategy that might be taken to deal with a novel problem. "Rule-based" action involves the application of a known strategy to solve the problem (either because the problem itself is familiar or because the problem is sufficiently analogous a familiar problem to allow rules to be transferred). Again, this approach was considered in chapter 2. Both of these approaches fall squarely within the remit of information-processing psychology. "Skill-based" action, on the other hand, owes a debt to Rasmussen's readings of Gibson's perception-action coupling. The second clause of the last sentence is deliberately clumsy because the relation between "skill-based" action and Gibsonian concepts is not as clear cut as some writers have claimed. However, the key point here is that the decision ladder represents a marriage between information-processing and "embodied" theories in ways that I have been trying to prevent.

Before exploring what an "embodied" reading of Rasmussen's thoughts might look like (that is, how one might read the word "ecological" in terms of radical embodied cognitive science [RECS]), it is worth noting comparisons that have been drawn between the decision ladder and recognition-primed decision-making (RPD) as descriptions of "expert" performance (see chapter 3). By decomposing activity into task steps, the decision ladder is concerned with defining what tasks need to be done while implicitly indicating that these tasks will be performed by one or other entity in the system, while RPD does not draw so clear a distinction between tasks (nor does it imply that the "system" involves more than one entity).[14] Further, RPD (in its assumption that the "expert" experiences the situation in terms of schema drawn from prior experience) proposes that there are features in the environment to which the "expert" can respond, while the decision ladder proposes that the salience of the features will be derived by the "expert" during their analytical activity.[15] For our purposes, RPD *could* be framed in "embodied" terms (chapter 3) rather than relying on the information-processing foundations of a decision ladder.

The decision ladder could be related to the task-artifact cycle in terms of ways in which to discover new possibilities or constraints on the information available to the system (which is also a core objective of CWA more broadly). In contrast, RPD (particularly in the "embodied" form we are advocating) is concerned with appreciating the ways in which particular sets of features in the environment acquire salience for the expert in order to better understand how to abstract these sets of features, either to better understand the expert's tacit knowledge or to provide cues and guidance for decision-making that could echo expert performance. In this sense, the ambition would be to design user interfaces that can enable "skill-based" activity. Beyond this, however, is the argument that "skill" is a far broader concept than Rasmussen seems to accept.

Defining Information to Support "Skill-Based Activity"

From the previous discussion, we could say that "skill-based" action involves recognizing and responding to affordances within the environment. In this respect, the word "skill" is intended to capture the accumulated experience that allows activity to be "automatic," or performed with little or no conscious awareness. It is the contention of EID practitioners that traditional user interface designs are created on the premise that all information is of equal value and that the person is able to extract the appropriate information for whatever decision they are making. Given the complexity of process control rooms, this can be challenging. EID[16] responds to a task ecology through the question of what information is necessary and sufficient to enable a person to contribute to the system achieving its functional purpose. From the work domain analysis (figure 5.3), one can determine the relationship between functions at different levels, and from the decision-ladder (figure 5.4), one can determine the strategies that people are liable to follow to perform these tasks, and hence, what information is required to support these tasks. More than this, the abstraction-decomposition space indicates which "values and priorities" constrain system activity and so indicates where information needs to be provided to show the status of the system in terms of these constraints.

EID seeks to reveal the constraints and relations in complex systems. In this way, the task ecology in which a decision maker operates is made clear, so that effort can be directed to making decisions rather than collating information in order to determine system state. EID is intended to support the user in diagnosis of problems through means-end problem-solving strategies. Additionally, the visualization of the system's activity in terms of these constraints and relations means that the operator will be able to respond to a set of features presented as patterns (i.e., the creation of a Gestalt of related information and also in terms of the response that users are expected to perform), relying on a perception-action coupling between what is seen and what action can be performed. In a sense, EID has the aim of "keeping related things together."[17] For the experienced user, the patterns presented by the user interface can be responded to in an automatic, skill-based manner.

The primary goal of the designer, in this instance, is to create a user interface that supports "skill-based" action because that aligns with the perception-action coupling that underpins embodied cognition. In other words, a well-designed display ought to support "at-a-glance" reading and "intuitive" response (both of which are recognized as attributes of well-designed displays in the HCI literature). While this is a logical conclusion to be drawn from this distinction, I wonder whether it is entirely consistent with the broad aim of EID. That is, the knowledge- and rule-based approaches seem to me to rely on information-processing notions of cognition, while the skill-based approach does not. So, either we need to produce separate designs, drawing on distinct theoretical traditions, or we seek to reconcile these into a single approach. My preference is to see whether we can continue with the embodied cognition approach and ground EID in this.

Ecological Interface Design

The foundations of EID involve a mix of concepts from cognitive psychology (specifically from information processing and the study of fault-finding and problem-solving), Gibson's notion of ecology, systems engineering, and cybernetics. From the latter, Ashby's "Law of Requisite Variety" would mean that a complex system needs a complex controller, meaning that the combination of human operator, sensors, and automated analysis of data need to be sufficient to mirror the process. The purpose of the user interface would be to expose the human operator to the complexity of the process in such a way as to enable effective control actions to be performed. Further, because a controller requires a model of the system that it is controlling, the operator must be able to access a "model" of the process. From the information-processing approach, such a model would involve the human extracting information-as-content and constructing a "mental model" to make sense of the process. From an embodied perspective, the human operator is considered to be part of the system and, as a consequence, the essential concern is to define the constraints under which the system behaves. We noted, in chapter 4, that Gibson regarded such constraints in terms of the physical laws that governed processes and the conventions (i.e., social, legal, economic, and so on) by which activity is permitted. The abstraction-decomposition space represents these combinations of constraints. Consequently, user interface design should reflect the constraints in the work domain in which the system operates. Fundamentally, this criterion raises the questions of how these constraints can be presented in a relevant manner for the user and what information is required for such presentation.[18]

Thus, the primary difference between EID and "conventional" user interface design is that the latter presents information-as-content, allowing the user to apply different rules of interpretation to extract information into salient combinations, while the former presents information-as-context pertaining to specific purposes.

The essence of EID is the visual representation of control-relevant relations. This reflects the genesis of the approach in process-control domains. Broadly, the overarching aims of the approach are to (a) reflect the task ecology as experienced by the people who work in that domain and (b) present sufficient information to support activity at the different levels of the abstract-decomposition space. For Rasmussen,

> Any control action activated through a work station, serves to change the internal, causal or intentional constraints to let them bring system state to the intended target. The interface should then represent the actual state of affairs in the work space in a way comparable to a representation of the intended or the useful state defined by the current goal, together with the situation-dependent "affordances" i.e., the options available for action on the constraints of the internal processes as defined by the physical design or on intentionality as defined by policies, practices, or regulations.[19]

In order to achieve these aims, Vicente and Rasmussen proposed two design principles:

1. For "skill-based behaviour": To support interaction via time-space signals, the operator should be able to act directly on the display, and the structure of the displayed information should be isomorphic to the part-whole structure of movements;
2. For "rule-based behaviour": Provide a consistent one-to-one mapping between the work domain constraints and the cues or signs provided by the interface.[20]

These principles derive from the aims of EID and can be elaborated further into concepts that inform design decisions.[21] I will illustrate these concepts with examples from the DUal REservoir System Simulation (DURESS)[22] simulator which is one of the best-researched examples of EID. The "mimic concept" draws on the idea (common in process-control user interfaces) of presenting the layout of components in a plant in terms of their structural relations, for instance, in terms of diagrams that show how components connect to each other.

Where EID differs from "conventional" user interfaces is that, in addition to the physical relations between components, there is an emphasis on functional relations. In figure 5.5, a series of valves (indicated by the black

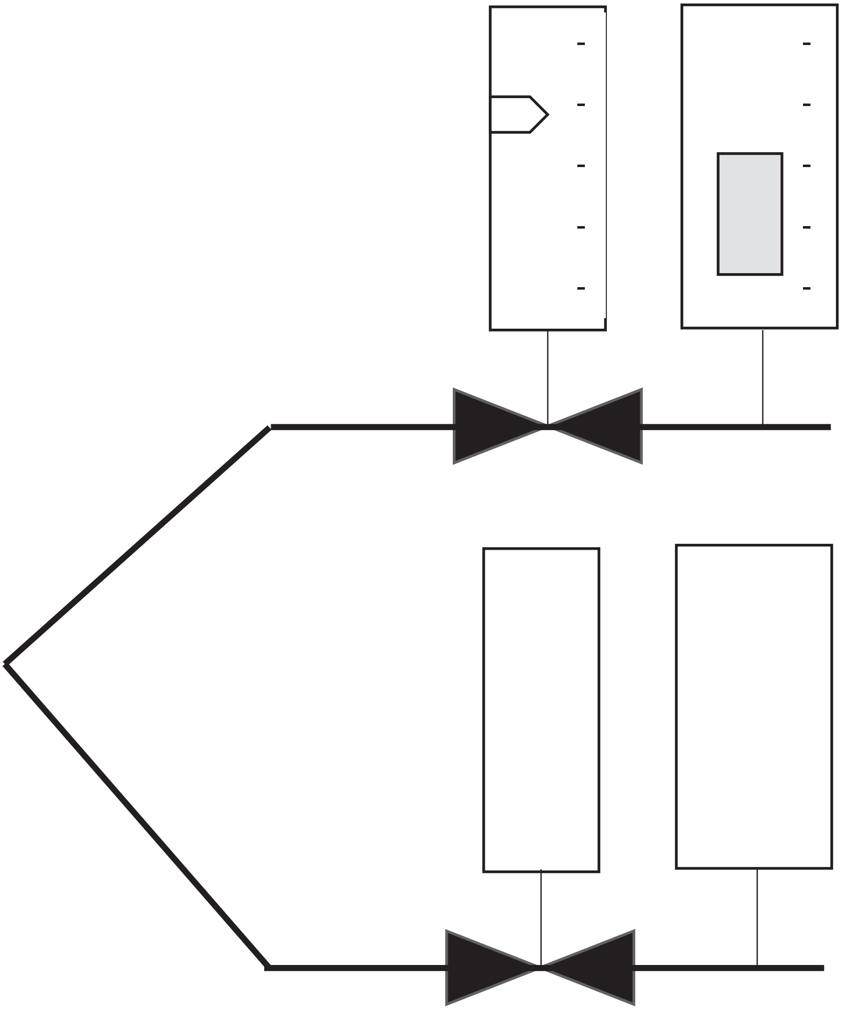

Figure 5.5
Indicating valve state and flow rates in an ecological interface.

"bow-ties") are connected to control the flow of liquid in a plant. That is, the user interface would show not just the "static" view of the system being controlled but how it changes state. In this way, the user interface will visualize the "flow geometries" that are relevant to the process (and are much like Gibson's "safe field of travel" for driving or walking, discussed in chapter 3). The "equality principle" states that some parameters describing the system ought to be equal under normal operations (and if these parameters deviate from these levels, then this indicates that the system is out of normal). In figure 5.5,[23] the status of each valve in indicated by a graph showing the flow rate (using a pointer on a scale), together with indicators of flow.

The "conservation principle" states that some parameters will be balanced under normal operations—that is, that energy lost in one part of the system ought to balance the energy gained in another part. One of the more commonly used examples in EID is the mass-energy balance display. In figure 5.6,[24] the quantity of material flowing into the process ("mass in") relates to the quantity of material flowing out. The values of these parameters are shown by the vertical bars, with a connecting line to draw attention to the

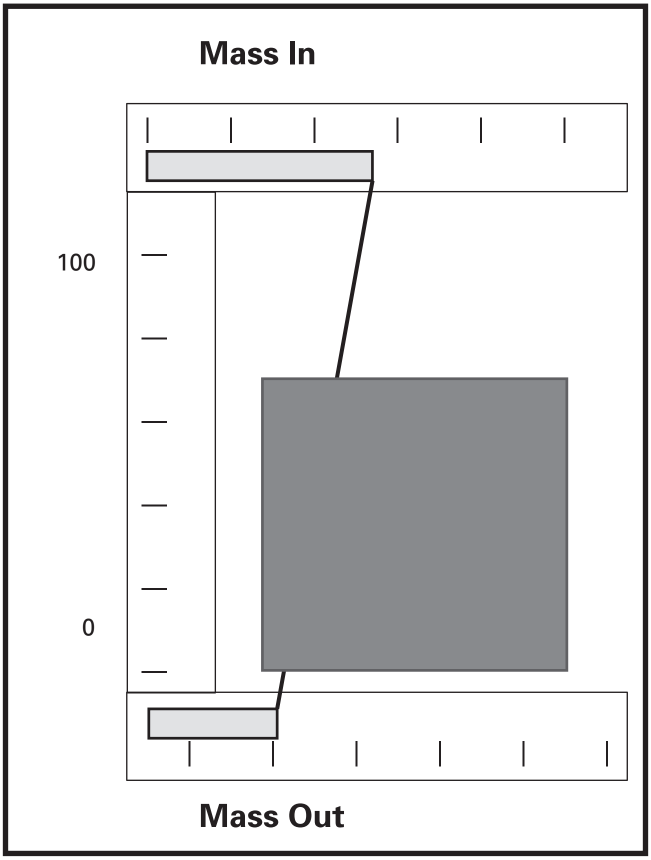

Figure 5.6
Mass-balance display in an ecological interface.

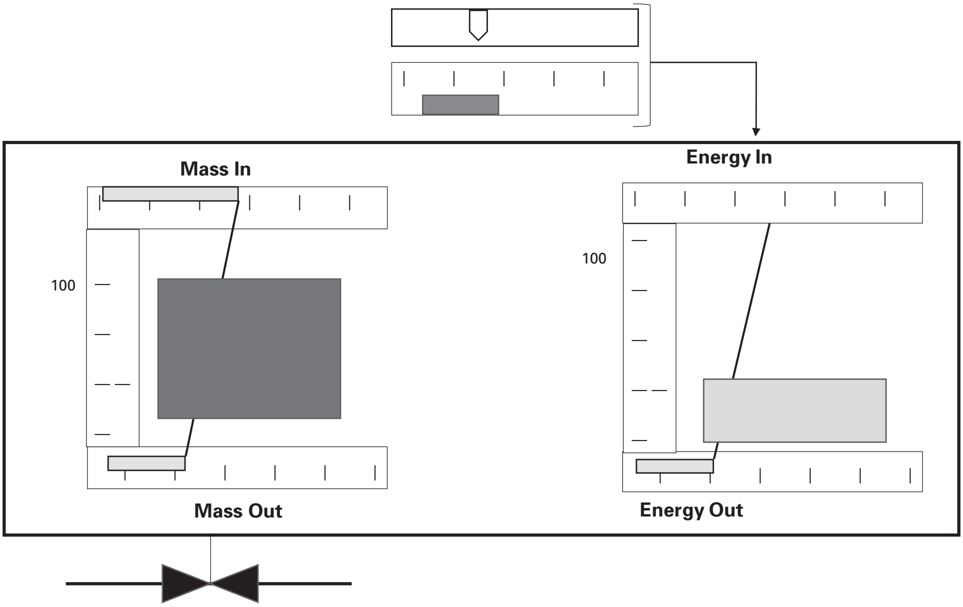

Figure 5.7
Graphically defining relations in an ecological interface.

relative difference within safe limits, defined by the square. In this way, the visual appearance of information creates a Gestalt, or pattern, that can be apprehended as indicating a particular system state (or deviation from that state).

Finally, the "pairing concept" states that certain parameters can be paired to produce a two-dimensional graph, as for temperature flow rate, for example (figure 5.7).[25]

Combining these different principles and concepts results in a composite display in which the elements provide different views of the "system" in terms of the abstraction-decomposition space. For example, figure 5.1 shows a concept user interface for military course of action planning.[26]

Are Ecological Interface Designs Better than Traditional Designs?

Gibson saw computer displays as providing "mediated indirect knowledge at second hand."[27] This would imply that their purpose could only be to display information-as-content, which the user would have to read and understand. The EID approach seeks to present information-as-context. In complex or nonroutine tasks, user performance (measured in diagnostic accuracy) tends to be superior for EID compared with "conventional" user interfaces[28] (although there is evidence that, even for normal operations, EID can be advantageous[29]). One explanation for this advantage is that EID presents information that is salient to understanding the function of the system in that situation while the "conventional" displays might require users to extract and combine information. The part-whole and means-ends approach to design means that the "window" of the process can be adapted to specific functions, which can enhance performance.[30] Indeed, adding information by moving up the abstraction-decomposition space, as the task becomes more complex, can enhance performance. In terms of decision time, some studies suggest that EID can be faster than "conventional" user interfaces, while other studies suggest that it can be slower. One suggestion is that this difference in time can be explained by differences in visual search patterns. However, some "experienced" operators are skeptical about the value of EID (even when using these designs produces measurable improvements in their performance) because they have a repertoire of interpretative strategies through which they define salience. When the "skill" of the user has developed to a level where they can "automatically" see patterns in their environment (or, at least, in the computer displays in their environment), they have developed perception-action coupling that allows them to determine the state of the process. The patterns presented by EID, in such cases, compete with these learned patterns. The question is whether the experienced users can, with practice, replace their old patterns with these new ones. One would expect the answer to be yes, but also that the "old" patterns, as manifestations of tacit knowledge, might be less easy to define that the EID patterns. That is, EID should allow users to "to better judge the validity and quality of specific computer advice,"[31]

What Does Ecological Interface Design Tell Us about Radical Embodied Cognitive Science (and Vice Versa)?

While the EID approach, with its acknowledgement of Gibson's contributions to "ecological psychology" and its version of "affordance," could share some common ancestry with RECS, I would like to argue for a stronger relationship. In part doing so calls for a separation of EID from the information-processing approach that CWA inherits from Rasmussen's early work on fault-finding in electronic circuits. In this early work (and the way in which the sequential decision-making seems to be mapped onto management of complex processes), there are some assumptions that seem to contradict, or reduce the importance of, the "ecological" stance that the approach offers. This is, I feel, compounded by the arbitrary division of "cognition" into "skill-," "rule-" and "knowledge-based behaviors" Vicente and Rasmussen noted that "because operators may engage in higher levels of cognitive control (e.g., knowledge-based action) even if the interface is designed to encourage lower levels, merely supporting the lower levels is not sufficient."[32] How are we to read this statement, if we want to argue against an "information-processing" approach (or, indeed, the notion of levels of cognitive control?). In the first place, this is advising against an overly simplistic assumption that there could be different forms of user interface that directly and uniquely correspond to the levels of control. That is, it does not make sense to design an ecological interface that is specifically for skill-based behavior and to have this as different in kind to one designed for rule-based or knowledge-based behaviors for the simple reason that people will approach the user interfaces with different task constraints and for different decision-making requirements. Consequently, defining the display specifically on the assumption of skill-based behavior could, paradoxically violate the broad aim of EID by presenting information that is either not salient to the task at hand or that requires interpretation in order to make sense of it.

In addition to the suggestion that it is a mistake to design EID explicitly for each "level," I am also proposing that the very idea of "levels" (skill-based, ruled-based, knowledge-based) is misleading and that it makes more sense to approach the design from the single perspective of RECS. First, both RECS and EID are concerned with understanding the relationships between constraints (actor, task, or environment) within an ecological niche. This means that the salience of information is defined by its utility in defining patterns in the environment that accord with actions that a given actor can perform. Different actors (defined by capability or by experience of that environment) could have different definitions of salience. Thus, it makes more sense to define information-as-context in terms of these definitions of salience than in terms of general notions of levels of cognitive control. That is, the EID for someone learning to control a process would differ from that required by an experienced operator—*not*, I think, because the "learner" operates on a knowledge- or rule-based level and the “experienced operator” on skill-based one, but because there are essential differences in kind between the ecological niches of these people. This relates to Brunswik’s notion of ecological validity in terms of salient information.

Both EID and RECS are concerned with the notion of “affordance.” As I argued in chapter 4, affordance is more than perception-action coupling (and definitely not some property of objects). As such, the notion is antithetical to an information-processing approach; you can have one or the other but not both because affordance emphasizes information-as-context and information-processing emphasizes information-as-content. However, the contrast between information-as-content and information-as-context raises a larger quandary in terms of what it is that the designer (of a user interface in this instance) is designing. If user interface design is a matter of providing information-as-content (on the assumption that users will apply their knowledge to interpreting, and discovering relations within, the information), then the key activities revolve around the definition of content. Indeed, in the mimic displays of process control rooms, each sensor (in the plant) was accorded its own “display” (often a dial on the wall-panel) positioned adjacent to the image of the element that it was sensor of. So, as in figure 5.7, each value would have a display to show the flow-rate through them (and the images of the valves themselves would probably be illuminated to indicate whether or not they were open). In this way, looking at the mimic display (or pacing along the length of the display, in a larger control room) would allow the operator to determine the state of the elements in the process. What this does not tell you, of course, is the relation between these elements (without significant effort on the part of the experienced operator).

6 Things That Think and Act

Introduction

Within the field of human-computer interaction (HCI) there are several ways in which technology is designed to act autonomously (or at least to give the impression that it is behaving in this way). This includes all manner of "smart" technologies such as robots and interactive toys, shape-changing artifacts, "things" in the Internet of Things, and software "agents." While this chapter is less concerned with the underlying algorithms that are designed to enable these various technologies to exhibit autonomy, the implications of how we interact with these technologies tells us much about what we might mean by "agency" (not only in terms of our interactions with "smart" technologies but also in terms of our interactions with other, everyday artifacts). Indeed, while it might seem obvious that, for instance, a robot has "agency" and a cup of coffee does not, I want to argue that (from the perspective of embodied cognition) this might not be as clear cut as we might assume. This is not because I want to claim that the cup of coffee has some malign intent that it is seeking to pursue (any more than I might claim that the robot is "evil"), but because "agency" and "intent" have to be considered in terms of the dynamic interactions within the human-artifact-environment system. The system, through the interactivity of its elements, seeks stability in certain states. From this, it becomes logical to assume that in certain states of such a system, the initiation of actions can come from either human or artifact or environment (rather than *all* actions arising from a single initiating, intentional agent) and that for a system "intention" could be equivalent to the states in which it is stable. This argument and its relation to concepts of agency are developed in the second half of this chapter. Before this, I want to give a broad sense of the ways in which "smart" digital technologies are developing. This is not intended to be a complete overview of the territory so much as a brief amble around places that I like.

Tangible User Interfaces

Tangible user interfaces have been designed with the aim of supporting "embodied interaction," which is "the creation, manipulation and sharing of meaning through engaged interaction with objects."[1] Unpacking this statement from Dourish, we can see three essential concepts that characterize his view of "embodied interaction": the first is "engaged interaction," the second is that this interaction is with "objects," and the third is that the focus is on "shared meaning." Each of these concepts, for Dourish, derives from his reading of phenomenology and, as such, they overlap with ideas in this book, but not fully. What is not so apparent from these terms is why they do not apply to *any* designed object. It might be the case that the emphasis on "shared meaning" implies focus on information-as-content—in other words, that the purpose of objects in a tangible user interface is to "create, manipulate, share" digital information. For the Tangible Media group at the Massachusetts Institute of Technology (MIT), led by Ishii, one can see that a logical extension of this ambition is to create physical objects that participate in the digital world: lift the stopper from a glass bottle and a sound plays[2] or move an object on a table and a projected display changes.[3] In these examples, action on the physical object mediates digital information. Rekimoto[4] presented a demonstration in which "digital objects" (files, movies, images, and the like) can be passed from one device to another, either by placing devices next to each other or by swiping the file to the top of the screen on one of the devices. Again, the idea is to allow actions to be performed on physical devices (connected to that same network and with digital objects identified by, for example, their web address). These examples have been foundational for HCI since the 1990s; while the underpinning technology has improved to create even more impressive demonstrations, the broad concepts share a similar goal.

Back in the 1960s, at the Stanford Research Institute, Doug Engelbart was leading a team of engineers who were exploring "next generation" computing. Inspired by ideas like those of Vannevar Bush and his concept of the "memex" (a glass table through which one could read microfiche and interact with the information these contained) and Ted Nelson's idea of hypertext, the work led to the "mother of all demos."[5] The oN-Line System (NLS) used a graphical user interface (far removed from the typical green or orange text displays of the time) to display information, represented by images of objects, such as files and folders, on the screen and offered videoconferencing and real-time collaborative editing of documents. Of particular interest for this chapter (apart from the fact that the research agenda of HCI was largely set by Engelbart's work for the next fifty years) is the way in which users interacted with the graphics on the screen. A wooden block with two wheels mounted orthogonally to each other (nicknamed the "mouse") was used to drive a cursor around the screen and buttons used to indicate a selection. Each aspect of NLS inspired research programs and the whole concept of contemporary HCI can be captured by the word "WIMP" (windows, icons, menus, pointing device).

While the focus of interaction with these devices remains the "information-as-content" displayed on the computer screen (and thus, these are not, by definition, forms of tangible HCI), the physical movements of, and with, these devices align with the points that I have been making in the earlier chapters and provide the basis for introducing radical embodied cognitive science (RECS) to HCI. It is not obvious that the design of tangible user interface is informed by theory of human activity. In his acceptance speech for the Association for Computing Machinery (ACM) Lifetime Achievement award, Ishii said, "There is no road laid out before me. I charge forward, and a road emerges behind me."[6] This echoes Varela's notion that enactivism is a path laid by walking, and so one might expect a theory of tangible user interface design to make reference to enactivism or embodied cognition. And yet, the theory presented to explain these designs makes little reference to these ideas. In part this might relate to the focus of the design activity. While tangible user interfaces produce ingenious and compelling demonstrations, the focus on information-as-content means that the devices become simply a different means of interacting with digital content. What is less apparent is how they could support information-as-context.

Anticipating embodiment in HCI, Winograd and Flores[7] drew on Heidegger's notion of "thrownness," which they presented as the experience of coping with the flow of interactivity between people and their technologies.

These aspects (experience, coping, flow, interactivity) are integral to the ideas of embodied cognition. Winograd and Flores discussed the ways in which technologies can become "ready-to-hand," by which they mean that an appropriately designed tool or piece of technology will "disappear" (from conscious awareness) during its use to perform a task. A similar notion is espoused by Marc Weiser, one of the early pioneers of "ubiquitous computing," who begins his best-known paper with the observation that "the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.[8]"

The concept of "ready-to-hand" also contributed to Dreyfus's'[9] critique of information processing (both in cognitive science and in artificial intelligence) because it emphasizes the importance of know-how in the grounding of knowledge (as opposed to knowledge being solely a matter of the manipulation of symbols). We become aware of the technology when the interaction breaks down (and the technology becomes "present-at-hand," i.e., merely an object that is demanding our attention).[10] Implicit in this observation is the assumption that break-down relates to know-how such that the more experienced user of a device is less likely to encounter break-down than the novice, for example, as a result of anticipation and skillful coping with the changing state of the artifact. The implication is that know-how increases our repertoire of coping strategies. This question of know-how and adaptation aligns with the discussion of affordance and of the skilled intentionality framework (SIF) in chapter 4 and with the notion of the human-artifact-environment system throughout this book. As we saw in chapter 4, one approach to capturing know-how has been to use concepts from phenomenology and practices from ethnography to produce rich pictures of the context of use and how this is experienced by users of technology. Dourish's own work has focused on the design and deployment of social computing, often through the use of tangible media.[11] Other writers have focused on the concept of embodiment as body-based interaction, influencing affect by encouraging different postures,[12] for example, or using degree of body movement to increase engagement in video-games,[13] or to interact with auditory displays.[14] Designs that violate "embodied metaphors" are regarded as less intuitive and harder to use than ones that follow these.[15] These body-based or embodied metaphor studies relate to Gallagher's description of minimal embodiment. Other researchers have taken a view that is more deeply influenced by phenomenology, particularly that of Merleau-Ponty.

In a study of a wearable device, worn on the forearm to support industrial maintenance, the device was interacted with by pointing the device at objects in the environment and tilting and tapping the display. Fallman[16] studied the skillful coping of users of this device in the workplace through interviews to elucidate user experience of the device in the physical and social contexts in which it was used. This approach captures the ecological niche in which the devices are used and aspects of micro-materiality (to use Heath's phrase from chapter 3) and illustrates how the environment (social and physical) creates the "landscape of affordance" that SIF describes, and captures some (but not all) aspects of the interactions in the human-artifact-environment system that concerns me in this book.

Somewhat closer to my aim is Hornecker's description of the ways in which "embodied facilitation" arises from, and creates, the structure in which actions are performed. For Hornecker,[17] the environment (whether it is realized in software or physical terms) facilitates activity through the ways in which features can serves as resources for action. To my mind, this description comes close to the system-level perspective on affordance presented in chapter 4. One can appreciate that "embodied facilitation" owes a debt to Merleau-Ponty's concept of intentional objects and the ways in which people respond to the opportunities and constraints that arise from the interactions between person, artifact ,and environment. While the trends considered so far have considered embodiment from the bodily, metaphorical, or phenomenological perspectives, there has been very little to date that has followed from "enactive" approaches,[18] nor for that matter, from the RECS approach followed in this book. In HCI, enactive user interfaces[19] provide closed-loop control by which the user directly interacts with (digital) objects through perception-action coupling—for example, through the application of notions of affordance to physical interaction with ecological interfaces (described in chapter 5). Consequently, while the attention that HCI has given to "embodiment" provides rich pictures of the context in which interactivity occurs, it does not (in my opinion) reflect the ongoing, reciprocal engagement that concerns me in this book. For this, we need to incorporate RECS into the design, evaluation, and study of HCI.

Tangible user interfaces allow people to hold and move physical objects to interact with digital information. Taking this a little further, HCI designs have explored ways of making these objects change shape. A simple example of this from my own work is the handle of a kettle which rises when it can to be picked up,[20] because the water in the kettle has boiled, or because the kettle in empty, or because the person should be using the kettle for the next step in making a pot of tea. We found that participants in our studies would, without prompting, refer to the kettle, jug, or other artifacts as "wanting" or "needing" that person to act. A more sophisticated example is MIT's inFORM,[21] in which hundreds of tiny motors drive blocks up and down to create a surface that changes shape. How these "shape-shifting" user interfaces "decide" to modify themselves leads me into a discussion of autonomy and "smart" technology.

Autonomy and "Smart" Technology

Technology that could sense and respond to its environment was the hallmark of cybernetics systems. While cybernetic systems operated with analogue data and contemporary systems operate with digital data, the basic principle of both is to sense (some aspect of) the environment, compare the sense data with a "goal," and act in order to achieve this goal. For this chapter, it doesn't matter whether the goal is to hit a target, avoid an obstacle, or retrieve some data: differences in behavior will arise from the degree of sophistication with which the algorithms cope with variation in the environment and the complexity of the goals that can be managed. How the devices relate data from their sensors to their action can be considered in terms of the "information-processing" and "embodied cognition" distinction I am making for cognition. An example of the "information-processing" approach is what Brooks[22] dubbed "sense-model-plan-act" (smpa), which follows the staged process that the phrase suggests. Of the several problems (for robotics) that this smpa approach creates are the adequate definition of the "model" that could be constructed from the sensor data and the definition and selection of appropriate "plans" that relate to this model. While computing power has increased considerably since Brooks was writing back in the early 1990s, these "deliberative planning" approaches still produce slow, hesitant robots. So, there are lots of short-cuts that engineers take to speed up the ways in which each of these stages can be performed. For Brooks, the solution to the smpa problem was to move from "deliberative planning" to "reactive planning" (although he was never enamored with this latter term) and to capitalize on four principles that define robots moving in their environments. These principles, inspired by biology, psychology, neuroscience, cybernetics, computer science, robotics, and the work of Agre and Chapman[23] reflect broad concerns of embodied cognition. Briefly, they are

- Situatedness: "The world is its own best model."[24]
- Embodiment: "The world grounds regress"[25]; that is, without the "world," then knowledge would be a matter of infinite regression where symbols are interpreted by symbols (with similar regression of homunculi to "read" the symbols).
- Intelligence: "Intelligence is determined by the dynamics of interaction with the world."[26]
- Emergence: "Intelligence is in the eye of the observer."[27]

While the first three of these properties can be read in terms of the discussions in earlier parts of this book, the fourth one suggests a mischievous turn. If you think about the automata of the eighteenth century or the robot "turtles" of Grey Walter in the 1940s, then you can appreciate Brooks's point. The former included clockwork "robots" that could write poetry, draw pictures, or play the piano, and the latter included small, three-wheeled robots that moved around their pen looking for "food" (their charging station) when they were "hungry" (battery power was low) and "playing" with (following or avoiding) each other. In these cases, the illusion of "intelligence" or "agency" was often compelling. In the first case, this was because the precise engineering of the automata's movement meant that they could reproduce actions that looked human-like (both to their eighteenth-century audience and to contemporary viewers who have the chance to see them in action). In the second, this was because the robots were able to sense and respond to their environment, and since the aspects of the environment that they could sense would alter by their movement, they looked to be adaptively responding.

The Furby (figure 6.1) was an interactive toy from Tiger Electronics, which first appeared in the late 1990s (with a couple of revivals from Hasbro in 2005 and 2012).[28] Furbies are animatronic toys that can move their ears, eyes, and mouth in response to sounds or touch or infrared signals (later versions replaced the infrared with Bluetooth). In this manner, Furbies can respond to being stroked or spoken to or to the presence of another Furby. Their behavior and appearance encouraged particular types of play (e.g., stroking, petting, talking to them) and, as they were played with, so

Figure 6.1
A Furby toy.

their "language abilities" and responses changed. They were programmed with their own "furbish" language and, with continued interaction with their owners, switched to a predefined set of English words. As their vocabulary had been designed to support particular types of interactive play, such as "how are you?," "tell me a joke," "tickle my tummy," they would switch from furbish to English on the basis of these types of play. This gave the illusion that they were "learning" English—to the extent that US National Security Agency decided to ban them because they might be robot spies. Whether this latter is apocryphal (my suspicion is that the fears were based not on language-processing abilities but on whether a Furby could act as a recording device—although, of course, whether "spooks" have time to play with toys is another matter) is less important than the idea that these toys were capable of intelligent interaction with their owners. Indeed, the more recent incarnation, the Furby Connect, incorporated Bluetooth connectivity through which it is trivial to hack into the microphone and speaker (so rendering this potential spyware or allowing someone to speak to the child who owns the toy—making this a far more sinister proposition and, perhaps, justifying the fears of the National Security Agency).

Being Digital

The purpose of digital technology is to digitize information so that it can be processed, stored, and retrieved from storage. Users of the digital technology issue requests (queries) to call up the information. These queries could take the form of verbal requests (typed, spoken, written) or physical actions (tapping, swiping, sketching). As of 2010, such requests can be anticipated by the technology so that we no longer need to make well-defined queries to get useful information. In part these anticipations are based on algorithmic models of similar queries or on the structure of the storage and connectivity of the information. And as of 2015, the technology is making inferences about our needs and intentions. We are living through a shift from a technology that reacts to our requests to one that anticipates our "needs." Such anticipations are based on models drawn from our prior behavior and a specification of what we might need. If the former clause of this sentence might feel neutral (although it is based on particular assumptions about how to model human behavior), the latter feels highly loaded: the very idea of technology satisfying "needs" is difficult to separate from the idea of technology (or the organizations that own, sell, or distribute the technology) creating or imposing "needs" on its users.

A *first-order* intentional system has beliefs and desires (etc.) but no beliefs and desires *about* beliefs and desires. A *second-order* intentional system is more sophisticated; it has beliefs and desires (and no doubt other intentional states) about beliefs and desires (and other intentional states)—both those of others and its own.[29] From this position, a thermostat is a first-order intentional system. It is capable of sensing its environment (i.e., having a "belief") and interpreting what it senses in order to perform an action (i.e., having a "desire"). Of course, these capabilities are so limited that talk of "intentionality" feels absurd. And yet, the question of what would need to be added to the capabilities for intentionality to be plausible is tricky to answer. For Dennett, the shift to second order intentionality involves the additional capability of having beliefs and desires about beliefs and desire—that is, meta-cognition, which allows the device to reason about how it might update its beliefs or about how it might achieve its desires. So, a "smart" thermostat could seek to balance the goals of making occupants of a house comfortable, minimizing energy consumption, minimizing energy costs, and so on. This could involve placing sensors around the house, access to information about energy pricing, models of the preferences and activities of the people in the house, and complicated algorithms that find optimal solutions to the competing goals. In this instance, the device is performing in a more sophisticated manner than the "dumb" thermostat. At what point (in terms of sensing the environment, range of responses, acceptability of response, and so on) would a device become "smart"? From a cybernetic perspective, the "law of requisite variety" states that a system ought to be designed to have at least an equivalent number of responses to the number of demands made on it (ideally, of course, each unique demand should have an appropriate response, so this is more than a count of responses and needs to consider what makes a response "appropriate"). In such a system, providing that the response matches the demand and the system moves to a state that is acceptable, the "intelligence" comes from its definition of environmental states and corresponding (or appropriate) action.

Negroponte, back in 1990, wrote an influential book called *Being Digital*,[30] which was a sort of manifesto for MIT's Media Lab and provided foundational concepts for digital technology. For Negroponte, much of our everyday life is (was) spent engaging in "analogue" activities, by which he meant a continuous stream of physical actions. The challenge was to convert these actions into a digital form that could be processed by computers. An example of this (and one that has became something of a running theme through the development of ubiquitous and pervasive computing) involves a refrigerator that is able to determine the volume of milk (or the freshness of the milk) in a bottle on one of its shelves and send a message to your cell-phone when you need to buy more milk (possibly sending the message to coincide with you passing a grocery store). The broad concept (over and above the algorithms and devices) is that digital technology becomes, in the words of Marc Weiser, woven "into the fabric of everyday life until they disappear."[31] For Weiser, examples of such devices ranged from tabs (such as badges that are networked to not only identify an individual but locate that person in their workplace and arrange for doors to open for them—assuming they have appropriate access privileges—and messages to be forwarded to devices local to that person), to pads (much like tablets with which we are familiar), to boards (public information displays). Central to Weiser's vision is the networking that would allow data to flow seamlessly between these devices. For some critics of Weiser's ideas, it is an oddity that he was not thinking of the cell-phone as the medium to support this vision, but I feel that this misses the point of his emphasis on data networks (which we have in many forms), on displays at different scales (which we also have in many forms), and on the seamless transmission of data between any form of display (which we do not currently have in the form he envisaged).

Dourish identified the importance of these shifting patterns of interaction to how we respond to and use digital technology with his call to the HCI community to focus its attention on “where the action is.” While swiping left or right to select “dates” might feel different from selecting an item from a list of options, or even typing a set of attributes and having these matched by an algorithm, the underlying cognitive processes involve expressing preferences and making a choice. A question is whether these preference and choice processes are equivalent. The information-processing approach, in which abstract concepts are manipulated in a mental model, could be interpreted as claiming that there is little difference between these activities in terms of making a choice (and there would be formal descriptions of decision-making that could support this). We saw, in chapter 2, that the way one expresses a problem and the way that one interacts with the problem space can have a bearing on decision strategy. So, does swiping a touch-screen constitute cognitive activity or is it merely a physical action that arises from (internal) cognition?

As a simple example, imagine walking up to a building in an unfamiliar part of town and taking your cell-phone out of your pocket. Having determined your location (through global positioning satellite data and a map of the town) and having run models on prior behavior, the screen on the cell-phone offers you the option to message the person you will be meeting in this building; perhaps the option would be to send the message “Hi, I’m outside. Can you meet me in the lobby?,” which would save you the need to type the message. Or perhaps the option would be to zoom in on a map to show that the building you need is a block away or to show an image of the building you should be looking for. In any of these instances, your need to search for information (or enter information) is reduced, and the technology is providing an invitation to act (e.g., agree to send the message, walk to the next block, look for the other building, and so on). Depending on your feelings toward technology, you might find this notion (or shifting the choice of action from you to your device) attractive or frightening. What is happening here is that the artifact is structuring and constraining affordances in the human-artifact-environment system.

Of the many questions that this example raises, some of the most significant center on the question of agency: Are you ceding agency to your device in accepting its recommendation? What can you do if the recommendations are inappropriate or unacceptable? What can you do if the underlying models that the device is using are inaccurate? Are you responsible for any consequences arising from following the device's recommendations? If the cell-phone's recommendation is not acceptable, what actions can we make? We could reject the current recommendation or request alternative options, or we could "retrain" the underlying model (or hope that rejection or selection of an alternative might result in the model being recalculated or including an exception based on this situation). Each of these accepts the underlying paradigm that the device is capable of making a recommendation based on our behavior and that we are content to receive and follow such recommendations. More than this, the fact that the interaction places us in a role as arbiter (allowing rejection, selection of alternative, or the possibility of ignoring the recommendation and doing something different) implies that any agency in this interaction lies squarely with us.

The "Internet of Things"

In the nascent idea of the "Internet of Things" (IoT) physical objects are equipped with sensors, processors, and communication capabilities to allow them to be networked. In a commonplace use case, a parcel containing items you have ordered from an online store can be tracked throughout its journey from warehouse to your hands. Tracking could involve reading a bar-code attached to the parcel each time it is passed from one person to another. In this way, the system keeps track of all the parcels it is managing (and it can readily answer the question "Where's my stuff?"). This is nothing more than a logistics supply-chain imbued with the ability to track items. The concept of IoT is intended to expand on this by allowing the elements (parcels in this case) to contribute to local (as opposed to centralized) decision-making. Perhaps the parcels could inform a sorting machine about their destination or delivery time; perhaps the parcels could inform the delivery person where they should be left. While there continues to be speculation about how IoT could be realized, this example highlights some of the challenges, particularly in terms of whether the parcel is represented as a physical object or as its digital counterpart. After all, it makes more sense for the details of the delivery of the parcel to be part of the logistics planning system than to be "known" by the parcel. As an "internet," the various sensors that respond to the parcel become connected and share information. This could allow the "parcel" (in either its physical or digital version) the ability to make decisions; that is, the parcel could define its own goals, seeking to optimize actions in terms of the information that is available to it (e.g., selecting a delivery time and location based on updated information from the purchaser rather than following the "standard" route of the delivery van). Rather than using a parcel-tracking application, I could have used sensors for traffic or pollution monitoring (where each sensor might adapt its data collection in response to the devices connected to it), but the basic questions about agency remain similar. Often in discussions of IoT, at least from an HCI perspective, writers confuse this concept with Negroponte's earlier ideas about a fridge that senses its contents and alerts you to pick up more milk. In this, the idea is that the "agency" is retained by the human (both in terms of receiving messages and deciding whether or not to buy milk—why buy milk if you are about to go on holiday?), whereas in IoT "agency" is retained by the technology (so, the fridge would place the order for the delivery of the milk and then either arrange for it to be delivered or create an errand for you to go and fetch the milk).

Switching our attention from physical "things" to software, "bots" are small pieces of software that can perform specific functions. For example, a bot could be tasked with seeking specific pieces of information on the World Wide Web (a "web crawler"). By and large, these are designed to sense a specific state of their "environment" and respond to this (e.g., by sending a report). These are "autonomous" in that, once launched, they will seek opportunities to perform the actions with which they have been programmed (not unlike a software version of Grey Water's tortoises) and that they respond to their environment (a malicious example of such bots are computer viruses).

In this section, I have skimmed over some examples of "smart" technology that can sense and respond to changes in their environment in ways that look appropriate. Each of these examples could be made a little more complicated by providing the artifacts with the ability to adapt their responses to different environmental states (i.e., to "learn" new relations between environment and response).

Levels of Automation

Rather than a binary distinction between whether technology is "smart" ("autonomous") or not, engineering uses the idea of levels of automation (LoA). In automotive engineering, there are five levels;[32] in ergonomics we tend to use ten levels.[33] In both schemes, the extreme cases involve situations in which the human is fully in control of an activity or the machine is fully in control. As table 6.1 shows, the intervening "levels" show differing degrees by which human or machine make decisions or work together.

In the LoAs 2–4 in table 6.1, the computer behaves as a "recommender" system. In LoAs 5 and 6, the computer makes a decision but seeks approval. In this case, the role of the human is not simply to approve the decision but also to fully understand and appreciate the consequences of the proposed action. This latter point makes more sense if you imagine that the computer is being used for safety-critical purposes. In LoAs 7–10, the computer decides on an action and then performs without allowing the human any opportunity to prevent its doing so. While these higher LoAs might sound unnerving, we have grown used to using technologies that follow one or

Table 6.1
Levels of automation scheme

Level of Automation	Description
1	The computer offers no assistance; the human must make all decisions and perform all actions.
2	The computer offers a complete set of decision/action alternatives, or
3	It narrows the selection down to a few, or
4	It suggests one alternative, and
5	It executes that suggestion if the human approves, or
6	It allows the human a restricted time to veto before automation execution, or
7	It executes automatically, then necessarily informs the human, and
8	It informs the human only if asked, or
9	It informs the human only if it, the computer, decides to.
10	The computer decides everything and acts autonomously, ignoring the human.

other of them. For example, many functions of modern automobiles, such as anti-lock braking, are applied in direct response to changes in environmental conditions and occur within a timeframe that is too fast to permit human intervention (although, in some automobiles we can turn this function off).

The Irony of Automation

While the idea of LoA has, for many years, informed design of complex digital systems, there are two fundamental problems that I have with this. The first is one that ergonomics has, since its inception, battled with. This is the very idea that the "human" always needs to be designed out of systems. Here, I am not making some Luddite argument against technology. In many instances, machines can do things faster, more accurately, more consistently than people. But these instances focus on the narrow concern of the activity itself and ignore ergonomics concerns or broader social and political implications. From an ergonomics perspective, "designing out humans" from systems creates what Lisanne Bainbridge called the "irony of automation."[34] There are three parts to this argument. First, in many instances, automation is not applied to every aspect of an activity. This means that automation is applied to those aspects that (technically or economically) can be automated, leaving a bunch of aspects that are not. This disparate collection of left-over tasks is then given to humans to perform in the service of the machine. Second, when automation goes wrong, it is the role of humans to intervene and put it right. This is where the "irony" starts to become apparent, because the disparate collection of tasks may not form a meaningful whole to the human, so it might not be easy to understand what has gone wrong. In order to support such understanding, a user interface displays information about the status of the automation, but the user interface provides a limited "window" on the state of the automation, and the information may require specialized knowledge to interpret. Third, having been removed from the "control loop" (through being designed out of the automated process), the human is expected to respond quickly, knowledgably, and correctly—and when this fails, the charge of "human error" is levied at the human who was unable to correct the failing. Designing humans out of automated processes is, of course, related to "deskilling" with its attendant implications for labor (both in terms of pay to workers and in terms of the ability of workers to define and protect their rights to recognize and preserve these skills through trade unions). If all work becomes deskilled, then anyone can perform it, and, if that is the case, then labor becomes replaceable and cheap. This is the argument against automation that has raged since Marx's critique of the industrial process and capitalism. For this book, there is a further aspect of deskilling that stems directly from the perspective of embodied cognition.

One of my favorite examples of a design that inadvertently designed humans out of a system concerns a large steel-rolling mill. In the old version, long bars of steel were heated and rolled to shape them; it was important to check the temperature of the steel, and the steelworkers could tell this from looking at the color of the heated steel. In the "new" version, the rolling mill was covered. Problems arose because steel was being rolled at the wrong temperature. So, temperature probes were placed inside the cover and operators were provided with a user interface, in which the temperature of the steel was converted to a color that was similar to that of the heated steel that the operators had gained years of experience in judging. In the old system, the ongoing, reciprocal engagement between steelworkers and the heated steel created opportunities not only for physical interaction (they would use poles to lever the steel as it moved on the rollers) but also to develop an understanding of the relationship between the temperature of the steel and its color. As any metallurgist or steelworker knows, this relationship is not trivial and depends, among many other factors, on the quality of the steel, the environmental temperature and air flow, and distance between furnaces. As this was truly tacit knowledge, it was not easy to put into words how the many factors interacted in this understanding. So, because the simple rubric that color equals temperature was used to design the user interface, the result was an unused and unusable display that could not support the steelworkers' knowledge or work practices. Ultimately, the new system was modified to include windows cut into the cover along the rolling track, so that the steelworkers could look in and see the state of steel. For me, a key point in this story is that the information-as-content (i.e., color equals temperature) did nothing to capture the actual knowledge of the steelworkers, and the information-as-context (i.e., the ways in which the color of the steel changed in response to a combination of factors) was lost. By mediating between the skilled person and the industrial process, technology had distorted and removed most of what made the process meaningful to the person. In terms of the irony of automation, the role of the steelworker had changed from one of actively monitoring and interacting with a process to one of passively monitoring a display in order to predict or guess when to intervene.

The redesign disrupted the ecological niche of the experienced steelworkers and replaced the affordances that were meaningful in their human-artifact-environment system with new affordances in a new human-artifact-environment system. The new system (with its color-coded visual display) removed them from the old system and lost the affordances that were meaningful to them. The solution was a clumsy attempt to refashion the original human-artifact-environment system (by cutting some holes in the guards over the rolling steel). In this example, the human had not been deliberately removed (designed out) of the process. Indeed, you might imagine that the design team had sought to do all they could to make sure that the role of the human was well catered to. Also, there was nothing that looks like intentional effort to deskill. However, the new design changed the task-artifact cycle in such a way as to alter the human-artifact-environment system and change the ways in which affordances arise.

From the "system" perspective advocated in this book, could we decide how "information-processing" gets shared between elements in the system? This allocation of function problem (for ergonomics) indicates the challenge of deciding which actor (human or automaton) performs which function. From the 1950s, the allocation of function has often been considered in terms of HABA-MABA ("Humans Are Better At . . . / Machines Are Better At . . ."), which implies a clear-cut distinction between a set of functions that are best suited to humans (such as intuitive problem-solving or empathy) and a separate set of functions that are best suited to machines (like lifting heavy objects or performing millions of calculations). However, a little reflection tells us that there are many, many functions that do not fit into the neat demarcation between human and machine. In reality, what tends to happen is those functions that can be given to a machine (within budget and within the machine's ability) will be given to it, with everything left over being given to the human. This leads to the "irony of automation" (see above). It is essential to design digital technology not as something that has humans as adjuncts but in terms of synergy.[35] A comparable argument has been made from the actor-network theory perspective.[36] Latour speaks of the "folding[37]" of human and artifacts such that they create mutually sustaining relations. An implication of this “folding” is that assigning “agency” solely to humans in these interactions does not always make sense. For this chapter, we note that actor-network theory relies on an ontology in which humans and artifacts are inseparable. From this, an “affordance” (chapter 4) is an instance of the folding of human and artifact in specific environments—that is, an affordance is the possibility for use of an artifact and this possibility is realized through the interaction between artifact and human in an environment. As a consequence, it is more useful to think of affordances not in terms of their meaning but in terms of their action potential (or their “behavioral meanings”)—in other words, to focus on information-as-context rather than on information-as-content.

Agency and Artifacts

In some instances, the software agent might have a physical manifestation. This could take the form of an intelligent digital assistant that responds to our spoken requests by playing music, giving spoken responses, or managing other devices in our environment (such as changing the room temperature or lighting, opening curtains, changing music volume, and the like). In related research, “chatbots” allow us to converse (either through typing or speaking) to a computer agent (typically in a well-defined domain, such as learning about geography or math). Similarly, virtual avatars present an animated character on the screen that responds to our questions not only through spoken response but also with changing facial expressions. More advanced versions of this concept have the “face” projected onto a fully articulated robot. We might ascribe “human-like” abilities to these technologies (in terms of the ways in which they simulate human behavior), but when they behave in ways that are slightly different from our expectations, we encounter what has been termed an “uncanny valley.”[38] In some instances this is simply a mismatch in terms of expected and actual performance, as in misunderstanding a question and supplying an erroneous response or presenting facial expressions that seem inappropriate to the context of the conversation. These effects can be quite subtle but will be sufficient to shift our interpretation of the “human-like” nature of the behavior; to use Heidegger’s phrase, the uncanny valley shifts these “virtual agents” from being “ready-to-hand”[39] to being “present-at-hand.” But there is a deeper sense in which the uncanny valley can be unsettling. This is not only where we are "creeped out" by the behavior but also where we realize that the "agency" we have been attributing to these agents is not as complete as assumed. That is, we might have assumed that our intelligent digital assistant or the onscreen avatar was capable of anticipating what we had been requesting. For example, we might assume that (like humans) these "agents" could respond to the illocutionary force of a comment as well as direct requests. So, you might say out loud, "I wonder whether mum has gone to the garden center?" to your partner, and the "agent" might overhear and initiate a phone call, saying "Calling mum . . ."

One reason why we might ascribe agency to artifacts (beyond the fact that the design of digital artifacts like the ones in the previous paragraphs are meant to simulate this) relates to the media equation. This suggests that we have a tendency to anthropomorphize many of things in our environment, from our pets to our automobiles to photocopiers that don't print when we want them to. But this also means that the notion of "agency" is more than simply the initiator of an action; rather, it becomes a question of how "cause-effect" relations (between an action and its outcomes) fit into the context in which these occur. As Latour observes, "The prime mover of an action becomes a new, distributed, and nested set of practices whose sum may be possible to add up but only if we respect the mediating role of all the actants mobilized in the series."[40]

I like the suggestion from Andrew Pickering that there is a dance of agency between user and artifact because it helps clarify the idea that in these interactions there is a loosely coupled "system" that is dynamically changing.

In Pickering's account of a scientist who conducted an experiment with a piece of equipment (a bubble chamber, perhaps like the one show in figure 6.2), there was a continuous series of moves by the scientist, who "sometimes . . . acted as a classical human agent; then he would become passive and the apparatus took over the active role, doing its thing."[41] What this illustrates is the performative aspect of doing (in this instance it is doing a science experiment, but the point holds across any domain). We perform an action on an artifact and it responds; how it responds then influences the next action that we can perform. But how its responds is also influenced by its properties (the material from which it is made) and its own environment (the various forces acting upon it). So, in a very real sense, our interaction is only partly about responding to the artifact and equally about

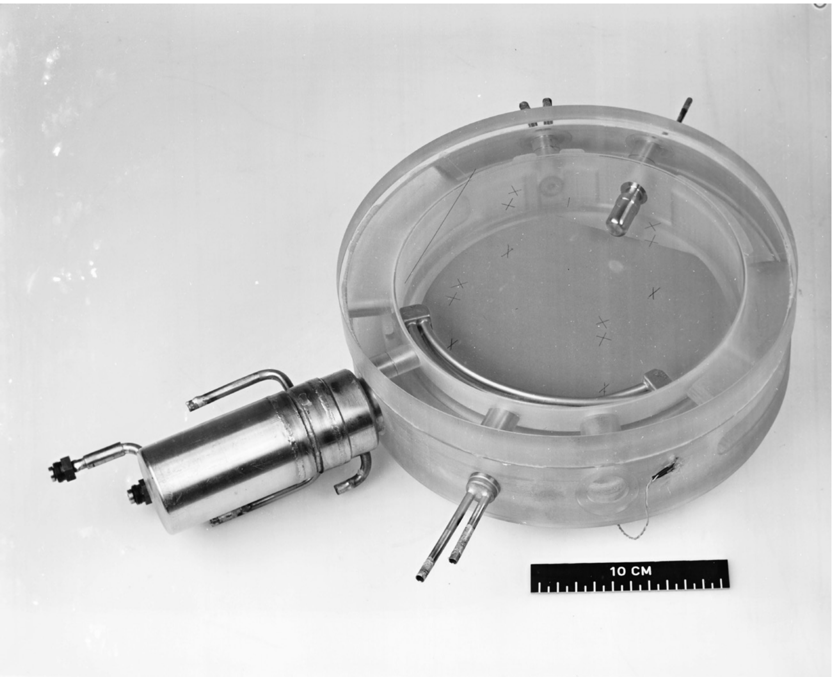

Figure 6.2
Bubble chamber.[42]

managing the behavior of that artifact in its context, or responding to it "doing its thing." In this case, the artifact is, in Maturana's terms "structurally determined":

> If you push a button something happens—it washes, it glows, it plays music—which is not determined by your pushing the button, but, rather, is triggered by the pushing of the button. . . . You do not instruct a system, you do not specify what has to happen in the system. If you start a tape recorder, you do not instruct it. You trigger it.[43]

Agency, Responsibility, and Theories of Mind

For many social scientists, "agency" evokes a capacity to act. Applied to artifacts, this requires either an internal drive (perhaps a motor, perhaps a processor that responds to different input) or an external force (perhaps as a physical force like gravity or perhaps a person, or animal, to act on it).

So, when you pick up a cup you provide the external force and when you drop it, gravity provides another force. To speak of the cup having "agency" might feel a little far-fetched (because one might assume that a capacity to act also involves sentience and intelligence and responsibility). So, how can there be a "dance of agency" (with its implication of reciprocity between human and artifact)?

How you reach for the cup, how you lift it to drink from, whether you perform another action (say, spill some of the contents to make it less full, or wait for it to cool down)are not simply matters of "making a decision" (in the sense of conducting some conscious calculation of risk to benefit). Rather, you respond to the artifact in context, perhaps while simultaneously having a conversation with someone. More significantly, how you shape your hand in reaching for the cup and how you move your hand toward the cup already contain the decision that you've made. As your hand reaches out, it is oriented to support performance of a particular action. These "reach-to-grasp" movements are considered further in chapter 7. However, these actions do not seem to involve thinking in a information-processing terms; rather, they involve thinking in the physical sense of interactivity that is the root of the ideas in this book.

In terms of the relationship between technology and cognition, one can posit a "weak" and "strong" view of distributed cognition. A "weak" view might claim that what is being distributed is the collection of artifacts upon which the act of cognition can be focused. This would require artifacts to play a passive role in the process of cognition and for them to function as vehicles for the storage or representation of information. Thus, the design of artifacts that are used in a work environment becomes changed by their use, and these changes provide cues for subsequent use.[44] The artifacts allow users to off-load information[45] and also provide a record of previous activity. In this version, the objects have their states altered by the actions that their users perform on them, for example, through note-taking, folding, or other markings. A "strong" view of embodiment might posit that it is the tasks involved in cognition that are being distributed. In order to accept the "strong" view, one must accept that "cognition" take place outside the head. If this is true, then many programmable artifacts (whether physical devices such as calculators or software "agents") can be claimed to be capable of cognition, as we saw from Dennett's arguments earlier in this chapter.

7 Recognizing Activity and Intent

Introduction

As shown in chapter 6, concepts derived from a broad appreciation of “embodiment” have made inroads into the discipline of human-computer interaction (HCI). The ways in which the various flavors of “embodiment” have been applied in this field are instructive, and they echo work across other design domains. Broadly, “embodiment” in HCI relates to the manner in which humans engage in physical interaction with artifacts and their physical surroundings. At one level, this concerns the design and development of tangible user interfaces discussed in chapter 6. As I noted, there is less work that has an obvious link to radical embodied cognitive science (RECS), particularly in terms of understanding the dynamics of continual reciprocal engagement within the human-artifact-environment system. RECS offers the potential to provide an account for ontology and epistemology in HCI (and interactions with other artifacts).

Within a human-artifact-environment, the environment can be characterized as a set of features, and a subset of these features will be salient to an individual’s task ecology. In the information-processing approach, salience relates to information-as-content (where features are selected on the basis of their meaning). In contrast, perception-action coupling emphasizes information-as-context. We have considered how Newell’s notion of organism, task, and environment constraints (discussed in chapter 3) might influence this notion of salience, and RECS can provide the framework within which to consider how salience is defined within the human-artifact-environment system. Moreover, salience relates to the objective that the system is optimizing (or satisficing). From an information-processing perspective, the "objective" could be a predefined "goal" that the human (as "intentional agent") is seeking to achieve. In terms of computer recognition of activity or intent, we might define the objective in terms of parameters for a Markov decision process, or "reward function" in reinforcement learning, or "priors" in a Bayesian model. In these instances, the computer will be seeking to optimize these values, possibly in terms of encouraging or discouraging particular actions. We discussed, in chapter 6, the notion of "agency" in interaction between human and technology and suggested that either the human is in charge, as the "intentional agent," or the computer is in charge, optimizing specific parameters. However, this contrast relies on an assumption that there is a directed relationship between human and computer which is contrary to the system perspective taken in this book.

Within the human-artifact-environment system, certain interactions will be possible (between human-artifact, human-environment, artifact-environment), and the outcome of these will create "states" in which the system is stable. These stable states represent objectives for the system (in dynamic systems terms, these are "attractor states" in which the system is most likely to rest). In this way, activity involves seeking a stable state (as opposed to seeking a defined goal). Of course, the stable state might correspond to a goal but the point of this description is that (just as we saw in chapter 3, Suchman suggested that plans can be situated, i.e., discovered opportunistically in the ongoing, reciprocal interaction), so dynamic systems show how stable states can be "discovered" through activity. The transition from one state to another will be determined by the constraints that apply to elements in the system and their interactions. RECS provides the framework for the ways in which these constraints affect activity. From the proposal that a designer should understand the ongoing, reciprocal engagement with the environment and that the environment offers a "landscape of affordances," an account that follows RECS should be able to reflect the richness and complexity of such interactivity. To begin the discussion, I consider the simple activity associated with reaching for artifacts to pick them up.

Reaching for Artifacts

People respond to affordances "automatically," with little conscious awareness of the features to which they are attending.[1] So, how is information acquired and used, in relation to affordance? In a series of neat experiments, Tucker and Ellis presented people with images of handled artifacts (say, a saucepan) and asked them to press a button with their left or right hand, depending on whether the image was inverted or not.[2] Irrespective of decision (inverted or not), response times were much faster when the handle of the artifact pointed toward the hand required for the response. The implication was that the orientation of the handle "primed" (that is, preactivated) movement of the hand that would be used to grasp the handle, and subsequent studies replicated this effect.[3] This suggests that the presence of an artifact that *could* be grasped initiates activity that would support grasping. It is a moot point as to how fine-tuned this relationship might be.

Reach-to-grasp movements are adapted to artifact properties. That is, the manner in which reaching is performed is affected by the artifact's width, weight, and slipperiness and by the subsequent action. The distance between thumb and fingers changes depending on the type of artifact that we will grasp. What is apparent from many of these studies is that there is a commitment to a specific form of grasp (as evidenced by the orientation of the hand and by the width between thumb and fingers as the hand approaches the artifact). A series of experiments have shown that it is possible to adjust this commitment during the performance of the action. This adjustment could involve the person reaching for an artifact and avoiding a set of distractor artifacts,[4] which slows movement time, or cuing a person to reach for an artifact in a specific location, and then changing the cue during their movement toward that artifact so that they need to change movement to a new location.[5] In anticipation of subsequent action, people adopt uncomfortable postures because they are seeking comfort in an end state (wine glasses) or in order to exert maximal torque (faucets).[6] Thus, the type of grasp to make when reaching for an artifact adapts to the relative weight given to artifact size, temperature, handle orientation, distance from the person, and so on. One implication (which reiterates points raised earlier in the book) is that we ought to attend to the most salient (rather than all) the available features, and that selection of features could be influenced by action pattern (in a reciprocal manner to the action pattern being influenced by the attended features). This ability to adapt movement, posture, and grasp according to the properties of the artifact or the demands of the task suggests that there is a process by which the exploration (of available features) occurs in an optimal manner—it doesn't make sense to assume that every detail is extracted and processed prior to performing an action.

In the simple act of reaching for an artifact, "decisions" are made and amended rapidly. For some researchers, this suggests that we have sophisticated "feed forward" control systems, in which a model of the world is used to plan and then guide our movements in the world. Other researchers take a very different perspective, arguing that our actions are model-free and that we respond to opportunities offered by artifacts in the world. It is this latter perspective that is most aligned to the claims made in this book. But, as figure 7.1 indicates, it can be challenging for a computer (using cameras to provide data) to determine that a person is reaching for a jug; it can be even more challenging for the computer to predict *why* the person is reaching for the jug or how this jug will be used.

Not only can we respond to the affordances of physical artifacts in anticipation of acting on them, we can also respond to the affordance of moving artifacts. A well-known example of this is the "outfielder problem,"[7] which relates to the challenge of catching a ball hit into the air. In order to catch the ball, you need to position yourself at the most likely point

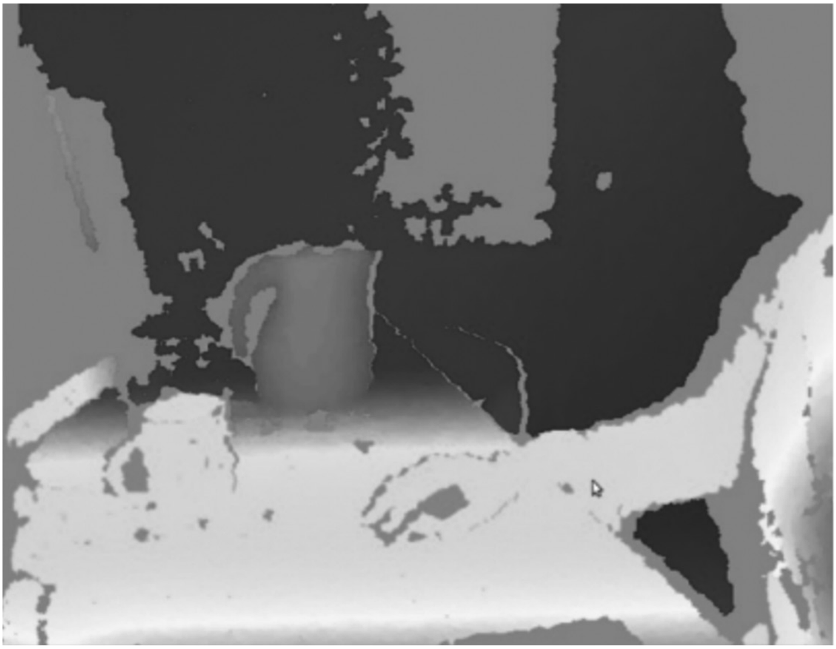

Figure 7.1
A computer's view of a person reaching for a jug.

to intercept it as it drop. The mathematics required to do this—direction, velocity, angle, and so on—are not beyond the ability of the average high school student,[8] but, as with many of the points in this book, the question is not whether are we capable of doing such calculations but whether we do these, as a matter of normal activity, or something else. There is plenty of evidence to suggest that, rather than performing complicated calculations, we employ strategies that rely on the visual relationship of the ball to catcher such that the ball appears to fly in a straight line[9] or to follow a constant velocity.[10] In either case, we are exploiting the optic flow in terms of a specific outcome in order to apply a simple (nonmathematical) adaptation to the visually perceived movement of the ball. These examples highlight the importance of considering the changing relationship between ball and catcher in the environment with the outcome being a "state" in which human (catcher)–artifacts (ball, catcher's mitt)–environment (sports field) be one of equilibrium—that is, ball in catcher's hand on the field. The "other states" of this system (e.g., the ball is dropped, the catcher runs into another player or advertising hoardings, the catcher trips, and the like) indicate the need to coordinate movement in order to achieve one outcome rather than others. The 'fly-ball' examples also illustrate the proposal that the "states" of the "system" can be important in defining the outcome (and provides a way of thinking about "intention" and "anticipation"). This implies that the action will be constrained by what constitutes a catch, which will be governed by the social conventions surrounding a particular sport, which constrain the types of actions that can be performed and the types of artifacts that can be used. One could imagine how different definitions of "catch" could result in different designs for gloves—for example, a "catch" might require the ball to firmly grasped by all fingers of the hand, in which case the glove would need to be fitted to allow each individual finger to move, or a "catch" might require the ball to be held in a container, in which case the glove would need to incorporate a receptacle that can hold the ball and need not separate the fingers. The contemporary baseball mitt is closer to the latter than the former but is still recognizable as a "glove" into which fingers are fitted. How these different glove designs define the concept of "catch" depends on the affordance offered by the context in which the human-artifact-environment system operates. That is, the human has capabilities (from prolonged experience of the sport) and wears the artifact (glove of particular design) in an environment (defined by both the physical ball park with the moving ball, other players, spectators, and so on, and the social conventions surrounding the sport, such as the rules) with the objective of having the ball end up in the glove.

The idea that we respond to patterns of environmental features is not simply to claim that animals (including humans) are "programmed" to respond to aspects of the environment around them. This would return us to "behaviorism" and its reductionist view of human ability. Rather, RECS recognizes that the salience of information defined by environmental features depends on the capability of the animal. For example, bees are able to see light in the ultraviolet spectrum, and flowers possess ultraviolet patterns. That is to say, flowers "mean" nectar to the bees, not because of some interpretative act but because of the complementarity between the capability of the bee and the property of the flower. While this might seem obvious, it leads to the radical proposal that the "interpretative act" that is the root of semiotics might be less about cognition (and information processing as interpretation) and more about complementarity. It might be, for instance, that seeing the artifact is sufficient to activate regions of the motor cortex associated with actions that *could* be performed with that artifact, and the immediate visual information from the environment is used to tune the subsequent action. From the perspective of enactive or embodied cognition, the "knowledge" that the person would draw on need not be represented in the form semantic knowledge of artifacts, but rather would take the form of a filter that actively selects "salient" information (however that is defined) from the environment.

From a RECS perspective, the relationship between a person's prior experience, the actions that they can perform, and the environment in which they perform these actions (which, of course, includes the artifacts being used) combine to form the perception-action coupling that underpins affordance. In other words, affordance arises out of an interacting web of relations. "Affordances are neither properties of the animal alone nor properties of the environment alone. Instead, they are relations between the abilities of an animal and some feature of a situation."[11] Affordances thus arise through relations in human-artifact-environment systems, rather than existing as discrete properties of any constituent component or as the outcome of some interpretive act. In other words, affordances represent stable states in the human-artifact-environment system, and, in order to define these stable states, we need to understand how the dynamics of this system are managed. RECS uses dynamic systems methods to do so.

In broad terms, the dynamic systems approach provides a mathematical basis for describing human activity. Much of the fundamental work in the application of dynamic systems to human activity is based on simple patterns of movement. This is for two reasons. First, the mathematical descriptions and analysis become increasingly challenging when we move to more complicated patterns of movement. Second, many of the underlying models can be characterized as coupled oscillators. The reason for the latter is both empirical (certain types of movement and certain types of neural activity can be described in terms of time-varying activity, e.g., rising and falling in signal strength) and theoretical (oscillators tend to have nonlinear relations to each other such that their dependencies are definable but not predictable). As an example, a classic demonstration of dynamic systems involves the simple action of tapping your index fingers on a table.[12] Begin by slowly alternating between right and left fingers to tap in sequence, and then speed up to as fast as you can. For most people, at some point the two fingers more together. That is, you begin by deliberately tapping out of phase but as your speed up, tapping becomes in phase. In terms of dynamics, this “system” has an order parameter that is defined by phase. Thus, the “system” has a tendency toward order in two states: fingers moving out of phase or fingers moving in phase. An interesting finding (typical of nonlinear dynamics) is that the transition from in-phase to out-of-phase is abrupt. That is, the transition between the two states of the order parameter is not gradual but happens suddenly. To appreciate these order and control parameters, we need to look at the basics of control theory.

Control Theory and the Human-Artifact-Environment System

If we begin with the view that activity occurs with the human-artifact-environment system, the question is how do the various types of interaction produce outcomes and how are the interactions performed as efficiently as possible? One way of thinking about this, as we discussed in chapter 1, is in terms of coordinating the interactions so as to minimize the degrees of freedom (DoF) of the system. Accordingly, coordination can be thought of as the process of selecting as few parameters as necessary to manage. From the theory of control systems (derived from cybernetics), we can say that the system defines an order parameter (which defines how well ordered, or stable, the system is in any given state) and a control parameter (which indicates the actions that can be performed to alter the order parameter). In this case, an order parameter could be a single variable or it could be the product of two or more variables, depending on how the system operates. In other words, the order parameter can be thought of in terms of the objective that the system is seeking to optimize (in line with our discussion in chapter 1). In the simplest version of this, a feedback loop compares the current state of the system with a defined level of the order parameter and calculates an error, which is corrected by performing an action. In this case, the aim of the system is to maintain stability (or homeostasis) in the face of disturbances arising from sources in the environment external to it. Systems of this type align nicely with the homeostatic models that Craik[13] proposed and that we considered in chapter 1. Rather than sampling the environment to create a model of it, the system defines a model based on a setting of the order parameter and its actions, are directed toward keeping the state of the system within the limits set by this model. Recall that the term "cybernetics" refers to a person who steers a ship, and you can readily see how such a model can describe a simple process in which deviation from a defined course is minimized.

From dynamic systems research, control-theoretic models can describe how DoFs can be managed. One implication for RECS is that the solution to the DoF problem requires the definition of an appropriate order parameter. A point to note is that the state of the system (perhaps obtained through sense data) is *not* a mental model of the environment to be constructed but an error to be corrected either to maintain the current state or to ensure transition to another state. In biomechanics, this concept can be considered from two perspectives. In one, movement is considered in terms of the "product of force (kinetic) fields and flow (informational) fields" such that coordination emerges from dynamic environments.[14] From this perspective, systems self-organize as they adapt to changes in the balance between internal (individual) and external (environmental) constraints. In the other, synergetics[15] can be described in terms of variability and consistency in movement, as reflected by "dynamical equations."

Given the range of actions that people could perform to achieve a specific goal, there is a Degrees of Freedom (DoF) problem (outlined in chapter 1).

An elegant solution to DoF is proposed by Bril,[16] who follows Bernstein in proposing a hierarchical control model. In Bril's approach, functional (order) parameters can be achieved through regulatory (control) parameters through which the person controls specific movement parameters. For instance, experts (flint knappers and stone bead makers) seek to hold the functional parameter (kinetic energy) constant when they use different types of hammer or material, while novices vary kinetic energy with different types of hammer. Recently, we applied this finding to the comparison of jewelers performing simple sawing tasks, showing how experience relates to the grip force applied to the handle and to the velocity of the saw blade during cutting.[17] As Bernstein noted, it is important to incorporate feedback into the closed loop control of motion, in terms of the interaction between person and environment. This feedback can be seen as a means of managing the dynamics of the human-tool-environment system. Rather than considering movement as the enactment of a program or schema, an alternative view is to consider the control parameters that need to be optimized. Thus, an optimal control model would seek to determine the "cost function" that is being minimized while allowing the goal of the movement to be achieved. Bernstein spoke of coordinative structures in which combinations of muscle activation become associated with specific movements in levels of synergy.

From these basic control-theoretic or comparator models, we can draw several conclusions that inform RECS. The first is that there is no need for a central "controller" to manage interactions because feedback loops between the system components will allow the output of one component to affect another. At some point, discrepancies between elements in the system will decease, and the system will be in a state of equilibrium. That is, through these feedback loops, the system self-organizes. The feedback loops create a circular causality in which prior states of components lead to hysteresis (literally "history matters"), but once the system achieves equilibrium, this constrains the value of the order parameter and brings the other elements into defined interactions. As the feedback loops can change the state of the elements in the system, the initial state of the system is important in defining how interactions might develop. What is critical in this approach is that we are less concerned with discrete interactions between elements and more concerned with the overall activity in the system. In this respect, the objective can be defined by the order parameter. The set of possible states in which the system is stable defines a state space. Moving from one state to another involves a phase transition, typically in response to an external change and typically in a way that is abrupt. As control parameters change, so the system shifts from one state to another. As the order parameter changes, the probability of moving to (or away from) a specific state increases. This means that one can consider these states in terms of attractors and repellents that pull the system toward (or push the system away from) parts of the state space. If we think of speed-accuracy trade-offs in a reaction time experiment (where you emphasize speed of response or correctness of response), the order parameter (time to respond to a signal) depends on two control parameters. We might claim that a strategy (favoring speed or accuracy) involves a phase transition that emphasizes one control parameter over the over.

Kalman Filters

Control-theoretic models of human activity were the direct descendants of cybernetics, in that they were feedback loops in which a servomechanism corrected movement in response to deviations from a defined path. As an example, imagine steering a car on a winding country road. Assuming that there is no other traffic or other obstacles, you could perform this activity by looking at the road ahead and making small corrective adjustments to the steering wheel to keep the car in the center of the lane in which you are driving. If this was all that driving involved, a basic servomechanism would be sufficient. Of course, we do not believe that this model describes driving because we typically have to attend to more things than the empty road ahead of us. However, as an initial example, this gives a flavor of basic cybernetic, closed-loop control.

Among the problems with the simple closed-loop servomechanism presented here, one of the most pressing is the way in which it handles uncertainty in the input signal. What you would not want was to drive the car by swinging the steering wheel in response to any perceived change in the environment. Not only would such control by ineffective, it would also be really uncomfortable for you and your passengers. So, this requires a way of deciding whether or not to react to changes in the input. A common approach to modeling manual control (at least for this sort of "tracking" task) involves the use of Kalman filter. The purpose of the Kalman filter is to reduce uncertainty in the input signal. In this case, the "controller" samples the environment and issues a control signal to maintain the state of the system within acceptable limits. For driving, “acceptable limits” will be defined by the position of the vehicle in its lane; if the road curves, then the vehicle needs to turn to keep in the middle of its lane. Samples from the environment could be affected by uncertainty (perhaps it is twilight or foggy or raining, so the road ahead is not so easy to see clearly). The “controller” needs to decide the degree of confidence to give to each sample before it issues a control signal. If the input signal has high levels of uncertainty, and the controller responds to all samples with equal confidence, this could result in very jittery control. Consequently, the Kalman filter compares each sample with an expected signal. The expected signal reflects the average of prior samples (as the input signal) and the current control signal. The decision to change the control signal depends on the confidence given to the input or output signals, together with the rate at which the samples were obtained in order to define and correct “error” (between input and output).

What is particularly important for a Kalman filter control system is to have a continual stream of information on which to base its analysis. Indeed, if there is no new information (either because the input signal has stopped completely or because there is no change in the input signal), then the controller becomes very sluggish in its response (because it cannot detect the error signal that it requires or because any new information might require a large adjustment). Thus, one might expect that the brain (if it behaved like a Kalman filter) would be continually sampling the environment, adjust body posture or move sensory organs to provide an ongoing stream against which it could update and maintain its model. Without committing to any claim that the brain is a Kalman filter, it is worth noting that the saccadic movement of the eyes[18] or a phenomenon such as postural sway[19] indicate a continuously varying input.

For me, the Kalman filter, as an error-correcting servomechanism, provides a simple analogy for how the brain might be for “coping not copying” in its interactions with our environment. I use an analogous argument to explain how the concept of recognition-primed decision-making (central to naturalistic decision-making) could be described as a closed-loop control system, so that it did not need recourse to schema or mental models.[20] What these mechanisms suggest is that (certain) activity can be described in ways that allow accurate prediction and that have no need of a mental model of the environment. Rather, they rely on models that reflect the "relation-structure" that Craik described (see chapter 1). The idea that the environment provides an "input signal" to a controller is not so far removed from the perception-action coupling of Gibson. However, the analogy of a servo-mechanism might feel dangerously close to a totally mechanistic (or worse, behaviorist) account of human activity; it might be acceptable to think of machines or robots as behaving in this manner, but how well does this fit with human behavior (especially if we want to capture "cognition" and "creativity")? Equally, is there a risk of replacing one form of internal representation (mental models and the like) with another (probabilistic or other weighting of salience)? Before answering this, I want to pose a counter-question: If one accepted that the information-processing metaphor (with all of the attendant baggage that I challenged in chapter 1) could describe how the brain functions as a "copying" machine, why balk at the suggestion that servomechanisms can be provide a metaphor for the brain as a "coping" machine? Both approaches (information-processing and servomechanisms) are reductionistic, both are based on machines, and both are intended to guide thinking through metaphor, and yet, only (I suggest) the servomechanism and its related concepts provide opportunity to rigorously describe how activity might be coordinated. I say this because the information-processing approach has a tendency to reduce itself to a set of interconnected "boxes" (describing particular functions), in a "production line," with an over-reliance on assumptions about what "information" is being "processed." A Kalman filter replaces the "production line" with a neater system that adapts to changes in the environment without the need for multiple stages of translation of "information." What a Kalman filter does not tell us is how features in the environment have salience. Kalman filters assume that the input is in a defined format (which is why a "tracking" or "steering" task is a useful way to conceptualize it). Nor do they tell us how the output (the actions we perform) adapt to environmental, task, or human constraints. To consider these issues we need to be able to define salience of cues and to define how the system manages its objectives.

The Bayesian Brain

One could characterize information-processing approaches to cognition as saying that cognition *causes* action. Indeed, there is so little consideration of action in a conventional information-processing experiment that the response a participant makes is often reduced to pressing a button. However, as we saw earlier in this chapter, even the act of pressing a button can be primed by the action context in which it occurs (i.e., pressing a button that is on the same side as a handle on an image results in faster response). This suggests that cognition and action must be more closely intertwined than most theories of cognition assume. Early accounts of action assumed a close-loop relationship between the brain and the environment. William James, for instance, used ideomotor theory to account for the ways in which humans learn to control activity.[21] Babies kick and wriggle and through these seemingly random movements begin to sense differences in afferent information, which, in turn, become available to perception and result in association between a specific movement and specific neural patterns. This is the basis of Hebbian learning, in which neural pathways become entrained and reinforced through the practice of specific movements. When the specific movement is required at a later time, these pathways become reactivated (assuming that they have sufficient resting potential and that they continue to be primed). While this process might account for our ability to perform specific movements, it does not seem to say anything about cognition. However, this misses the point that such learning creates the ability to intentionally perform action—although, of course, there is still the feeling that the input to this intentional control comes from some “cognitive” activity.

One way of conceptualizing a closed-loop control for human action is to use the Bayesian brain approach,[22] which assumes that the brain does not act as a passive filter; rather it (1) has a set of probabilistic models, “Bayesian beliefs,” of the sources of information available to the senses, and (2) uses these “beliefs” to make predictions about how the information will change. As soon as there is a discrepancy between the prediction and the information, there needs to be either an effort to collect more information or to change the beliefs. There are some similarities between the manner in which the beliefs are used to define sources of information and a Kalman filter discussed previously and the manner in which these beliefs are updated and the Brunswik lens model discussed in chapter 3. To a great extent, the problems associated with the Brunswikian model apply to the Bayesian brain concept (i.e., an assumption that the world is sampled in terms of internal states and that these internal states are used to determine action). On the other hand, one could interpret the Bayesian brain in cybernetic terms as an example of Ashby's law of requisite variety (in that the Bayesian beliefs should be sufficiently complex to create expectations of the state of the environment relative to a person's actions). Indeed, in the Bayesian brain literature, there is assumed to be a "Markov blanket" in which a given state can be predicted because the model contains sufficient states to make such a prediction. A potential problem here lies in the scope of the blanket; as with the law of requisite variety, there is an implication that the "model" can contain all possible states that the system will encounter. In a cybernetic system, say, geared to managing temperature or water pressure, one can imagine that a finite set of states can be defined, which are sufficient to explain activity (and even here, one probably needs to have a couple of "wild-card" states to reflect unusual causes of puncture or damage to the pipes or damage to heating elements). But would one commit to the idea of a sufficient "model" for the brain in its interactions with the environment?

In a Bayesian description of brain activity, beliefs are specified in a hierarchy of layers in which high-level goals are defined as the prior probabilities, which then influence lower layers (Friston claims that this hierarchical structure can be found in the cortical structure and that it involves the activity of pyramidal cells).[23] Sensory information can be broadcast across the brain, and this requires adjustment (of the gain of channels over which the information flows) so that specific prediction errors can be managed. Consequently, optimizing the operation over the different layers can involve seeking and reconciling error between what is expected and what is observed. Such adjustment involves a process in which prediction errors are minimized by either updating the priors or by seeking additional information—that is, "active inference" is performed to guide sensory activity to reduce such errors. Consequently, this approach has also been termed "predictive processing."[24] The overarching goal of this activity is to maintain the brain in a state in which entropy is as low as possible (in other words to avoid increased entropy or "surprises" arising from uncertainty).[25] As soon as entropy increases, actions are performed to collect more sense data (or, by analogy with control theory, to modify the order parameter), which is used for comparison. Given that collecting data (or modifying the order parameter) can have associated costs, a further goal is to ensure that such costs are minimized.

This Bayesian brain approach describes some aspects of the ongoing reciprocal engagement between human and environment in terms of continuous perception-action cycles in which the order parameter of the system in defined in terms of sense data. As with our previous discussion of multiple objective optimization (in chapter 2), we can assume that there is a large number of states in which the brain can operate, but that its state at any given moment will be defined by a much smaller subset of states (with the aim of maintaining equilibrium or homeostasis as far as is practicable). If there are prediction errors (due to a mismatch in the current level of the order parameter or to sudden changes in the environment), these will increase entropy and cause homeostasis to be disrupted.

The Bayesian brain system seeks to manage "free energy," which, from information theory, means that the brain seeks to minimize any discrepancy between belief and sensory information in order to keep the long-run average surprise (unexpected or out-of-model) events as low as possible.[26] Free energy depends on incoming sensory signals, conditional expectations, and a model that relates conditional expectations to states of the world. This results in a scheme in which conditional expectations are replaced by sensory signals and the model is updated. The error (between prediction and model state) then informs the resulting response. From this perspective, perception is not a process of creating mental model that contributes to cognitive processing, but a means of managing sensory information within constraints set by the prior probabilities of information in the world that the brain is configured to respond to. Action, from this point of view, becomes a way of either updating these prior probabilities or seeking further sensory information. Accordingly, the purpose of cognition is to maintain homeostasis of activity in response to salient information.

The Bayesian Body

The Bayesian brain hypothesis provides a way of theorizing how affordance operates and an elegant set of testable hypotheses about how the choice between seeking further information or performing an action is made. At present, the Bayesian brain (and predictive processing) seems to situate all of the activity in the neural architecture of the brain and to rely on data from brain imaging to provide support to the argument. From an enactive and embodied perspective, this is troublesome because it offers little opportunity to include the body in the theory.[27] One approach would be to align a Bayesian brain approach with sensorimotor contingency theory.[28]

In this approach, the environment acts as the "external memory" from which to derive action. Sensory signals correspond to actions in this environment such that "rules or regularities relating sensory inputs to movement, changes and action."[29]

If one accepts that prior probabilities of perception-action pairings are adapted as features are attended to (and that some of these priors persist, perhaps as resting activations, between situations), then one could also accept that, by analogy, resting activations exist for the body. In a sense, this is what Bernstein meant by coordinative structure (see chapter 1). As an athlete or craftworker repeats a particular action, so the musculoskeletal system becomes tuned to that movement. In other words, "goals make perception enactive."[30]

As activity is performed, the interactions between elements in the coordinative structure vary, depending on the way in which the structure is being controlled and the way in which it is affected by the environment around it. This notion frames the point made by Ingold that there is a moment-by-moment, stroke-by-stroke variation in the tool-wielding movements of the skilled craft-worker.[31] In this way, one can consider activity in terms of softly assembled systems in which activity is contextually constrained and embodied and in which repetitive actions share a "family resemblance" but exhibit variability. Local interactions among embodied processes on different timescales weave the intrinsic fluctuations of the component processes into a coherent fabric of flux, despite inherent tendencies of the different processes to vary at their own different rates (on their own timescales). In other words, the challenge for understanding activity is less one of understanding discrete actions and more one of understanding the ways in which activity balances between consistency and variability, which is what Bernstein defined as "dexterity." In other words, skillful coping is not simply a matter of performing an action but rather is about acting in order that the human-artifact-environment system reaches a state that matches an objective, or an order parameter. As each element in the system might be subject to change, there is a need to adapt to change. Such adaptations, in dynamic systems terms, can be considered in terms of changes to the human-artifact-environment system, which can be measured in terms of stability or instability of the system. Measures of stability, over time, are derived from various definitions of entropy. When a system is stable, it will be low entropy. "Competitions among local rates of change strike a precise balance with globally emerging cooperative activity. In the precise balance of (or near) the critical state, they produce a long-range correlated, aperiodic pattern of change or flux in behaviour."[32]

With entropy analysis, we are in a better position to understand the underlying dynamics of activity. One approach is to use 1/frequency (1/f), which describes the fluctuations in time-varying data between highly predictable and totally random. In other words, it provides a measure of the underlying stability of the system that generated the signal. What makes this measure interesting is that many phenomena produce time-varying data that at local levels appear random, but that over longer timescales show repeatability. From this, 1/f scaling can also be considered in terms of long-term memory in signals. The reason why this is of interest to dynamic systems models, particularly in terms of human activity, is that it allows us to make sense of activity that might look unstructured or random on a moment-to-moment basis but that demonstrates a repeating pattern over many instances. To take a simple example, recall the reaction-time experiment in which you have to press a button each time a light turns on. Your time to respond ("reaction time") is a standard metric for a host of cognitive studies. Usually, the results of thousands of trials will be collected, and the average (mean) and variability (standard deviation) reported. What these statistics do not reflect is the way in which your attention (and enthusiasm) for the task might wax and wane, particularly over thousands of repetitions. If, instead of averaging reaction times, we treated these actions as a series of events over a time period, we can explore the strategy that is being applied.

Such 1/f scaling can be applied across different cognitive tasks to indicate a "softly assembled" system by focusing on interaction-dominant dynamics (in which component dynamics alter interactions) rather than component-dominant dynamics (in which behavior arises from components, demarcated and assigned specific functions).[33] In part, 1/f scaling reflects the motor component of the activity being studied and the ability of people to adapt to situational demands as embodied systems. For example, hand-mouse coordination in a simple video game exhibits 1/f scaling during normal operation but not when the task is disrupted.[34] This result indicates that during normal operation hand-mouse control can be described as an interaction-dominant system. Applying this concept to jewelers, 1/f scaling can distinguish skill levels in the use of jewelry saws.[35] In addition to dynamics being detected in physical performance, these are also apparent in cognitive and perceptual tasks. 1/f scaling has been shown in cognitive tasks,[36] and dynamic systems measures can be applied to reaction time experiments[37] and problem solving.[38]

Computer Recognition of Human Activity

The ability of computers to recognize and respond to human activity has grown dramatically since the 1990s. In this section, by “human activity” I mean speaking and moving. The proficiency of speech-recognition technology and wearable fitness monitors is such that these have now slipped over from being technology (with all its implications of the magical and beguiling) to the status of a commodity (so quotidian that we barely notice or question it operation—until it goes wrong). When I was doing my PhD on speech technology in the 1980s, the majority of speech-recognition algorithms (particularly for commercial applications) would use a limited number of words and a highly restricted syntax for combining these words, often requiring a period of “training” so that the device could modify its models to your manner of speaking. Most of these devices seemed to favor a sort of transatlantic English and struggled with pronunciation or accents that deviated too far from this.

Even with the major leaps in algorithmic complexity, both speech technology and wearable devices are essentially signal-processing devices. That is, the basic challenges in their operations arise from the collection, cleaning, and analysis of data from their sensors (microphones, inertial measurement units, and so on) so that these data can be used to create models against which new signals can be compared and classified

For the most part, speech- and activity-recognition technologies are concerned with isolating discrete “units” (e.g., words, actions) from the continuous stream of data coming from the sensors. One might assume that, for speech, the unit could be human-scaled—for example, a word. Unfortunately, defining such units at the “word” level tends to produce quite poor performance. This is partly because isolating words as discrete units can be challenging; in speech, words overlap and run into each other. This leads to problems of “end-point” detection (where each word begins or ends). The acoustic parameters of the signal can be affected by “linguistic” context (the words before or after it) or the “extra-linguistic” context (the emotional state or age of the speaker, the background noise). Noting that few speech-recognition systems make use of detailed semantic knowledge (they are, as we noted above, sophisticated signal-processing systems), these problems are not dealt with through understanding the meaning of the words. Rather, speech technology (since the 1990s) has focused on “units” that can be assumed to be fairly stable, or at least to have variability that can be predicted. To this end, the “units” are phonemes (or the acoustic equivalent: sounds that can be labeled as phonemes), with statistical models defining the probability of phonemes being combined in sequences. This is the basis of Markov models, which heralded a step change in speech recognition in the 1990s and was the basis of many commercial systems. Recognition performance (particularly in a benign environment of the laboratory) could reach above 90 percent in terms of accuracy—so you would have to repeat one or two words out of every ten. The advances in this technology over the intervening years have been remarkable, particularly with the widespread use of deep neural networks. In deep neural networks, the statistical patterns are discovered by computers through the correlations between phonemes in massive corpora of speech. For the purpose of this discussion, it is sufficient to accept that speech recognition involves the definition of discrete “units” (phonemes), that these units are probabilistically related to each other, and that all the information required can be obtained from the speech signal. Consequently, while the signal that this technology processes contains human speech, it is a moot point as to whether it “knows” that is dealing with “speech.” My point is that few, if any, of these technologies begin their analysis from an understanding of how a person produces speech.

The recognition of human activity (i.e., movement) can be performed either from sensors on the person or with cameras. For example, in the Microsoft Kinect a depth-camera captures the image of a person and this is translated to a point cloud that is matched to a skeleton model. As long as there is good alignment between point cloud and model, the person’s movement is recognized (so the avatar on the screen follows the movements of the player). However, the alignment might not be perfect, and players often need to subtly change the way that they move in order to maintain alignment. A similar adjustment happens with speech technology, such that speakers might alter their pronunciation or choice of words (particularly when the device has made a mistake). To date, much of the work using sensor data makes assumptions similar to those used in speech recognition (not least because so much of the analysis of sensor data either uses statistical models, such as hidden Markov models, or uses deep neural networks). While the issue of whether speech technology knows that it is processing "speech" (noted above) does not affect its overall performance, for activity recognition I think that are many unresolved issues. For instance, assuming that "actions" can be defined in terms of discrete units that are separable from the flow of activity is quite odd when applied to everyday settings. In some cases, say, counting steps on a digital pedometer, the model could be quite simple, as in defining a threshold for the signal to pass in order to count as a "step." Having said that, step-counting based solely on sensor data is not as trivial as this implies. In particular, deciding when a step has been completed could involve reconciling more than one impact (heel striking, knee locking, weight transfer on to front of foot, and so on), depending on the way that a person was walking (particularly if this person was relearning how to walk following an injury or was wearing braces or calipers), and on the location of the sensor (in the shoe, on the waist, in the pocket, and so on). Furthermore, counting steps is only part of the analysis that one might wish to make—for example, analysis of gait might be more important, particularly in rehabilitation.

For basic activity recognition, action can be defined in simple terms of a threshold beyond which the incoming signal needs to pass (as in the example of step-counting). A more complicated approach might combine parameters from several sensors to cope with contextual factors that could influence the signal. In this "context-aware computing" the challenge is to ensure that data from all the sensors can be combined into reliable models. In almost all cases, however, the models have little need to know about how the signal was produced. That is, these technologies rely on the assumption that all the necessary information can be extracted from the sensor data. The sensor data are then used to create a model. The model is used to evaluate any future sensor data, labeled using the "units" to which the model has been trained. Of course, this is the same process that the information-processing approach to human cognition adopts. My claim is that neither speech or activity recognition nor the information-processing approach to cognition begins its analysis from an understanding of how people produce speech or action. That is, rather than engaging with the embodied nature of human behavior, these approaches assume that this behavior can be reduced to a discretized model.

Some people who develop wearable technology or activity-recognition algorithms might be affronted by my claim that they ignore embodiment. I can imagine them saying something like, our devices attach to the body so we must be doing embodiment, but, perhaps, the majority will shrug and say, so what? Why should technologies that respond to human activity need a concept of embodiment? For wearable technologies, the concept of "activities of daily living" is commonly invoked to define classes of activity into which patterns of sensor activation can be grouped. One reason why these technologies might benefit from a concept of embodiment is to enable them to achieve their aims of adaptation and personalization. Rather than detecting *what* action has been performed, they could ask *how* it has been performed.

Many algorithms used for activity recognition are based on normalization of the data. That is, the models might identify statistical points of consistency, say, a central value in a cluster of similar data, and then create a boundary around this point to define an inclusion zone; any value within this zone would be treated as equivalent to the central point. So, if we give the central point a label, such as a particular phoneme or action, then any subsequent data that fall within the inclusion zone would be given the same label. By definition, this approach seeks to eliminate variability. In signal-processing terms, this makes sense because sources of variability might include noise or other interference to the sensor data and this needs to be minimized to reduce recognition error. But recall that Bernstein's definition of dexterity was based on adaptive variability in human actions. From this, it is unlikely that activity recognition, as it is currently performed, could adequately reflect the ways in which skills are learned (or lost). A common approach to "skill" (in activity recognition) is to define discrete levels, with a model defined from each level. While this aligns to signal-processing approaches, it does not align with theories of human performance or skill acquisition.

A little reflection on the suggestion that "skill" can be discretely compartmentalized shows its failings. Skilled practitioners do not always do different actions than novices do, nor do they always perform tasks more quickly. Rather, a characteristic of skill is the seamless merging of tasks into sequences and the ability to rapidly adapt performance of a task to suit context or the ability to anticipate the needs to a subsequent action and adjust a current action accordingly. Treating actions as discrete units misses that seamless merging (unless, of course, one creates models that reflect all possible combinations of sequence). If we do not use discrete units to define action, then we shouldwe treat actions as sequences, in time-series, which returns us to our earlier consideration of dynamic systems.

Recognizing Actions and Inferring Thoughts

We are all familiar with “recommenders” on websites, which suggest that, as we have purchased product X, we might want to consider product Y. Early instance of these recommender systems were based on crude matching of purchases (which could often lead to peculiar recommendations). Contemporary versions incorporate more nuanced reasoning and more information (often obtained through “scraping” the records of your interactions on a variety of webpages or with credit cards or store loyalty cards). In this case, the recommendations are developed from a detailed “model” of you as a consumer. While the idea of such models might be worrying (not least because we have little control over who is using our data and for what purposes), the point at issue here is how we are meant to respond to recommenders. For consumer decisions, these might be relatively benign (irritating but easy to ignore). However, there has been a growing class of recommender systems (often running on devices that we wear or carry) that are expressly designed to modify our behavior or change our habits. At present, these apps tend to be focused on health, particularly diet, exercise, smoking cessation, or medication reminders. These apps take data from sensors on the person (such as accelerometers or step-counters) and use these to provide motivational messages, or they have reminders programmed to occur at specific times, such as when to take medication. From the early 2000s, the input to the reminders comes from a broader range of sources; we have already considered location-based adaptation, for one, and personal information assistants can adapt to our previous actions and preferences.

Do I wish to claim that these devices somehow “know” what you are thinking? This sounds pretty far-fetched, particularly when you consider the type of data that such a device might be collecting. But is the idea of a wearable device that *is* able to know what you are thinking (or “read your mind”) simply a matter of the type of data that it collects? In a sense, this is only a matter of refining the ways in which “recommender” systems currently work. After all, if you plug enormous quantities of well-curated data into deep neural networks, then some consistent and intriguing results are inevitable. This is not a matter of opinion; it's just math—but, of course, it implies a particular definition of "what you are thinking."

Across much of cognitive science, "thinking" refers to purposeful, goal-oriented activity (such as the problem-solving we discussed in chapter 2), rather than the tumbling chaos of chatter that might intrude on our quieter, less distracted moments. In other words, "thinking" is typically defined in terms of a goal or intention toward which action is directed, rather than the mudging of thoughts about relationships, finances, or what to have for dinner. There is good reason for this focus, in the cognitive sciences at least. If you are going to study "thinking," then you need to know when it is happening, and, if you are running an experiment, you need to make sure that what is happening, happens in a similar manner to all the participants (otherwise you run the risk of the experiment being confounded by individual differences). In other words, thinking, for these experiments, involves the manipulation of information-as-content. Even conceding a narrow definition of "thinking," there remains a challenge of associating an action (or sequence of actions) with an intention and whether such an association necessitates a "theory of mind." For some sequences, this could be a trivial challenge. For example, you fill a kettle with water and put it on to boil, then you open a cupboard door and take out a cup. From knowing the time of day and detecting these actions (e.g., using data from sensors on the handle of the kettle and the cup, the door of the cupboard etc.), it would be probable that your actions will result in making yourself a cup of coffee—and your intention would be "make coffee." At this level, talk of "thinking" might feel redundant. More significantly, does the identification of a sequence of actions that can be associated with a known outcome actually signify intention? Before answering this, let's add a further element to the activity. Suppose that, as a New Year resolution or on medical advice, you have decided to reduce your caffeine intake. As long as one of the next actions in the sequence does *not* involve taking the coffee jar from the cupboard, then we can simply switch the notion of intention to "make a hot drink" (at a higher level of definition) or "make a herbal tea" (at a lower level). The suggestion of higher- and lower-level definitions implies a hierarchy of intentions (which can inform activity recognition by digital technology). But, let's assume that you have picked up the coffee jar. The sequence of actions now points to an intention of breaking your resolution or ignoring medical advice. In this case, the device that is monitoring your actions could intervene, perhaps by activating a buzzer on your wrist, perhaps by sending you a text message, perhaps by logging this intention, and sending a message to your physician. With this trivial change of context, this example has shifted to something that the reader might find more sinister. The shift has not come from a change in technology or algorithm (in each case, sensors generate data that are interpreted by algorithms tuned to detect and respond to specific features); rather it has come from the change of emphasis from recognizing activity to predicting intention. In the first example, the algorithm defines a "goal" (i.e., a class of activity that is specified by a collection of actions). Pursuit of the goal could be supported by, for example, having the cupboard door handle light up to cue the person to find the cup. In the second example, the algorithm is evaluating the activity in terms of a value structure, in which the "values" represent social or other forms of interpretation of acceptability of an action. This returns us to the discussion, in chapter 4, about the politics of affordance. In this case, the "affordances" relate to the opportunities for action that the device is defining for a given context and raises questions of how we, the users of the device, can accept or dispute such a definition and what options are available to us if we do not agree with the device.

If we are not directly interacting (or even not interacting at all) with smart technology, how should we consider our relationship with it? It feels as if some of the traditional views that HCI offers become redundant, as do the options that an information-processing approach might suggest. For example, if the behavior of the smart technology is opaque, should we simply seek to make it "transparent"? There is a lot of interest in the question of "explanation" of the decisions made by complex artificial intelligence. In this respect, the problem (of transparency or explanation) becomes a matter of information-as-content. My problem with this is that we are probably no longer interacting with smart technology in ways that make explanation possible or plausible. Any "feedback" that the technology presents to us will, at best, create further demands on our attention and decision-making and, at worse, become confusing, misleading, and pointless. However, if we think about how people provide explanations, we might realize that they are every bit a matter of information-as-context: not only do we adapt the content that we provide to our audience, but this adaptation often unfolds and develops in our conversation with them.

8 Eventually Everything Connects

Introduction

What is implicit in this chapter's title (taken from the quotation under figure 8.1[1]) is that there is great deal of trial and error in the continual engagement between designers and their materials before "eventually everything connects" and a design classic such as the Eames chair (figure 8.1)[2] is produced. The notion of continual engagement is at the heart of this book and underpins the concept of radical embodied cognitive science (RECS) it uses. By way of summary, I claim that it is important to adopt a systems view of interactions between human, artifact, and environment. This means that focusing exclusively on any one component or pairing will miss the subtleties and complexities of these interactions. It also means that claims about artifacts becoming "part" of the person (or incorporated into their "body-schema") miss the point. Within the human-artifact-environment system the boundaries between components remain demarcated, but the borders that allow exchange of information and action are permeable, and this is what gives rise to synergies and interactions. Before proceeding, I want to point out that there is little in the discussion so far that requires us to focus exclusively on "digital technology." All of the points that I have been making can apply equally to tools or to everyday artifacts such as kettles, jugs, and cups. When we pick up and use a physical artifact, we are part of a human-artifact-environment system. The artifact participates in interaction not only between it and the human, but also between it and the environment, and also mediates the interaction between human and environment. During activity, this mediated relationship provides information-in-context (in the form of feedback through various sensory channels) that

Figure 8.1
Eames Chair. Charles and Ray Eames were "able to make plywood bend to their will and yield the iconic Eames chair. Not in one week however. From alpha to omega, this project took years to culminate in the final chair. Countless sketches, revisions, practical tests were stepping-stones the Eames took in the process of their designing. Eventually everything connects."

can be used to regulate movement, maintain balance, correct for errors or deviations, and so on.

Creativity and Design

For Glaveanu, one of the most prescient writers on creativity in the 2000s, "Creative action is distributed between multiple actors, creations, places and times."[3] This idea implies the sort of "system" considered in this book and echoes an emphasis on the "organism-in-its-environment."[4] In terms of the human-artifact-environment system, much of the laboratory-based study of creativity has focused on the "human" and has had little concern for the interactions between human and artifacts (or the roles that environmental, physical, and social factors play). However, a review of problem-solving demonstrates how important it can be to allow the person to physically manipulate objects (and that the visual presentation of the problem can influence the choice of strategy to apply). The environment (for the artist, designer, or other creative practitioner) consists of the materials to be worked in order to produce the outcome and some cultural appreciation of what constitutes an acceptable outcome. From these, the creative practitioner works within constraints imposed by the nature of the materials, the types of tool being used, their experience and expertise in working with these, the design brief, and aesthetic, historical and cultural traditions and conventions. The creative practitioner will not satisfy all of these constraints but will work within the problem space defined in terms of one or two constraints. More broadly, one could see the "conversation" between design and these constraints as the ongoing (re)definition of objectives (both in terms of how these are evaluated but also in terms of which objectives to work with). None of the preceding discussion would come as a surprise to the creative practitioner, but I believe that very little of this fits an information-processing approach.

From a degrees-of-freedom perspective, the essential features of creativity are that there are some constraints that are fixed and a few that are open to adaptation by the practitioner. This suggests that some manipulations (of tools and materials) would be more likely to contribute to the outcome than others, either because the nature of the materials constrain some manipulations or because the outcome is defined by the sociocultural milieu in which the work is performed, which, together with the materials, tools, and other features of the environment, provide constraints in which creativity thrives. Indeed, creativity involves the cycling through acting and evaluating outcomes in such a way as to imply a closely coupled, dynamic feedback loop.[5]

As shown in chapter 5, the notion of "patterns" informs ecological interface design. Such patterns can also be explained by the phrase information-as-context, which I use in the book as a gloss on some of Gibson's ideas. Moreover, emphasis on pattern, symmetry, harmony, and balance echoes advice that can be found in any textbook on art, architecture, jewelry-making or any other form of creative practice. Knowing that certain patterns seem particularly salient to human perception, one can assume that "good" design (at least as far as one considers the form of an artifact) will carry with it some sense of the designer's interpretation of proportion, symmetry, harmony, or balance, each of which represents different objectives that the designer is using to constrain the design, In this way, the "outcome" of the design has, inherent in its definition, some or all of these concepts.

From this perspective, the role that a mental model has in design practice could be minimal. Rather, the physical interactions change the states of the human-artifact-environment in ways that create new opportunities for action and new affordances for activity to move to new states. For Sawyer and de Zutter,[6] creativity is opportunistic in that the creative practitioner responds to the "micro-affordances" that arise from their ongoing, reciprocal interaction with their environment (and the materials it contains). In this moment-by-moment shifting of contingency, potential synergies arise and provide constraints on potential actions. My intention is not to deny or denigrate people's experiences of creative practice. If the phenomenology of creativity involves a strong sense of ideas coming from inside the head (or from external muses or divine inspiration), I am not seeking to destroy such beliefs. Indeed, prodding such experiences might damage the very practice we wish to understand. As a surgeon once remarked to me when I was questioning how he made sense of his skill and practice, "There are better ways of understanding how a watch works than pulling it to pieces," especially (and here he looked over his half-moon spectacles at me) "if one is not a watch-maker." But, in my experience, creative practitioners are less likely to speak of ideas being fully formed in their heads and then made by their hands than to have a feeling in which a form reveals itself as material is being worked (or a form resisting particular demands of the maker). From an information-processing account, it is easy to dismiss the idea of a "form waiting to be found" as mysticism. Yet, this does seem to be a valid experience of many craftspeople.

As an example of this way of thinking, recall Karl Marx's contrast between spiders, bees, and humans:

> A spider conducts operations that resemble those of a weaver, and a bee puts to shame many an architect in the construction of her cells. But what distinguishes the worst architect from the best of bees is this, that the architect raises his structure in imagination before he erects it in reality. At the end of every labour-process, we get a result that already existed in the imagination of the labourer at its commencement. He not only effects a change of form in the material on which he works, but he also realises a purpose of his own that gives the law to his modus operandi, and to which he must subordinate his will. And this subordination is

> no mere momentary act. Besides the exertion of the bodily organs, the process demands that, during the whole operation, the workman's will be steadily in consonance with his purpose. This means close attention.[7]

The first part of the quotation echoes Aristotle's notion of "hylomorphism." Ingold[8] points out that typical accounts of creativity seem to assume a chain of causality from "idea" to "artifact." However, this is in stark contrast to the observations of artists and designers who might recognize the statement that "design is always a search for something that is unknown in advance."[9] In other words, creative practice more likely involves an interplay between forming (as a process of physical manipulation of materials) and thinking (as a process of evaluating and anticipating forms as they are made). Marx says that the "will" is in consonance with the purpose. In other words, the loose coupling within the human-artifact-environment system involves an ebb and flow of control, from person to environment and back. In this way, "creativity" could be a matter of sensitivity to changes in the way that we interact with the world around us. This is similar to Ingold's argument about creativity as improvisation.[10] In improvisation, there is a sense of foresight (rather than prediction) in which one is able to see, from a given point in time, a few steps into the future (but not to the very end of the creative process). This sense of improvisation would suggest that the designer responds to affordances in a current state as a means of spotting opportunities for action to transition to another state (e.g., through exploratory sketches, model-making, or working with different materials).

The Importance of Task Ecologies

To better understand how the human-artifact-environment system works, the notion of task ecology was discussed in chapter 3. In this, the ecological niche provides the resources for action, and here, the "ecological niche" consists not only of physical features but also of cultural conventions and practices. Good design is aware of these interconnecting features of the ecological niche. Indeed, I would go as far as to say that good design is concerned with the creation of ecological niches as much as it is with the creation of artifacts. While this claim might make sense when speaking of architecture (which, as a discipline, has a strong tradition of theory, much of which chimes with the points I am making in this book), I believe that it is also true of single artifacts. In essence, the design is constraining the objectives (of the human-artifact-environment system) in such a way that activities become apparent and that, for each of these activities, the artifact can be “ready-to-hand.” It is my contention that theories of design that see the artifact as the “solution” to a problem not only inherit some of the thinking from information processing that I have pointed out in this book, but also create overly constrained designs.

Even for a trivial information display, the notion of task ecology can be instructive. In an analogue clock (figure 8.2a), the numbers on the clock face and the moving hands indicate a specific form of information that supports a specific enquiry—namely, what is the time now? For other decisions, it is necessary to manipulate this information. So, the clock answers the question what is the time, but not questions such as can I make the next train? or how much longer is left for this lesson? The TimeTimer design, in contrast (figure 8.2b), supports the latter query by color-coding a region that indicates when the lesson in meant to finish (if lessons begin on the hour and end after 45 minutes, we add a red band on the clock face). It does not, of course, allow us to answer the question “What time is it now?”

The modification of the display is intended to support a specific enquiry—so the space of possibilities is open (for the clock) and closed (for the TimeTimer). As the environment becomes richer in information, the challenge of closing the space of possibilities for specific enquiries becomes greater, particularly when information sources increase, or some sources are more reliable than others, or when there are “incomplete invariants.”[11] Having said this, I would note that the TimeTimer is an example of a design in which the task ecology matches a specific goal in such a way as to support “direct perception”; in other words, the person is able to perform a specific task without manipulating the information. That is, it affords the task of deciding when a lesson is (or should be) ended—assuming that the person understands what the color indicates, that the clock time is correct, and so on.

While the TimeTimer introduces the idea of designing products to suit highly specific goals, I am not arguing for this as a general approach to design—not least because it encourages a commitment to “modes” (that is, highly constrained states of a product that are only appropriate for a single activity). While there might be places in which modes are beneficial, the general consensus in human-computer interaction (HCI), and design in general, is that these ought to be avoided because they limit the functionality and utility of a product. However, we can learn from this example that

Figure 8.2
(a) Analogue clock and (b) the TimeTimer countdown clock.

the space of possibilities offered by an artifact can provide constraints on action. The ecological interface design approach, considered in chapter 5, draws attention to the ways in which an understanding of constraints can benefit design thinking. These constraints, of course, apply only under certain conditions—for example, TimeTimer assumes an environment in which activity has a duration of forty-five minutes and that humans know the convention that the activity will end after this duration. What this artifact "affords" depends on the nature of the activities permitted within this environment and the conventions surrounding these. For example, is it acceptable for pupils in the classroom to start gathering their belongings together and standing up to leave when the hand has passed the forty-minute mark? Should the teacher beginning summarizing the lesson at thirty-five minutes? These questions also point to the reason why problem-solving is a poor metaphor for design; it implies that there is an impetus to produce a "solution" that can be expressed in a form. But a "poor" design may prematurely commit to a single state and constrain opportunities for new states to arise. This is another reason why the use of "modes" is frowned upon in design.

Affordances

A designer cannot simply imbue an artifact with a specific property and expect this artifact to possess an "affordance." Indeed, what does the idea of giving an artifact an affordance mean? I have argued against a notion of affordance as "artifact x affords action y." If affordances guide action, then this could only be the case for an agent able to perceive relevant "information," able to perform the relevant action, and able to relate the action to a desirable goal. In other words, design is not a simple matter of "fossilizing" a single affordance (i.e., a defined state of the human-artifact-environment system). If so, whose affordance and for what purpose? Alternatively, should design reflect as broad a range of potential affordances as possible, each arising from different states of the human-artifact-environment system (even ones that the designer has not yet considered)? In some instances, there may be overlap between these different states for some coherence and consistency to appear (or the ways that humans adapt their actions to the form of the artifact might allow consistent patterns of action to apply). What the designer is doing is not simply specifying a form but also laying out the affordances in which this artifact participates. As I noted in chapter 4, affordances are *not* merely properties of the artifact (if they were, then that would mean that an "affordance" is trivially the same as the "function" supported by the "form" of the artifact). Rather, affordances are the states in which the human-artifact-environment system finds stability. By defining these potential states, the designer is exploring the *some* of the range of uses, interactions, activities in which the artifact can participate.

The concept of affordance makes sense as emerging from a human-artifact-environment system in which the environment (and the artifacts it contains) combine with the human (and their capabilities) to create a set of constraints under which activity is performed. These constraints are further modified by the task constraints (relating to conventions surrounding the definition of a goal, criteria for the quality of performance, and so on) such that executability conditions influence the constraints within which the person can act. The action possibilities, in turn, become “intentions.” If we reverse this process, then the perception of features or performance of actions can create higher-level goals.

In order to explore the concept of affordance further, and to simplify the idea that there are different levels of “affordance,” I developed the idea of forms of engagement.[12] In this, the focus is on the ways in which we engage with artifacts and how different forms can serve to support and constrain each other. The most recent version of this concept is illustrated by figure 8.3. The arrows are intended to indicate the relation “constraints.” At the center of figure 8.3 is a dotted box labeled “affordance.” This describes a relationship between the ability to recognize salient features in an artifact (environmental engagement) and the ability to act using that artifact (motor engagement).

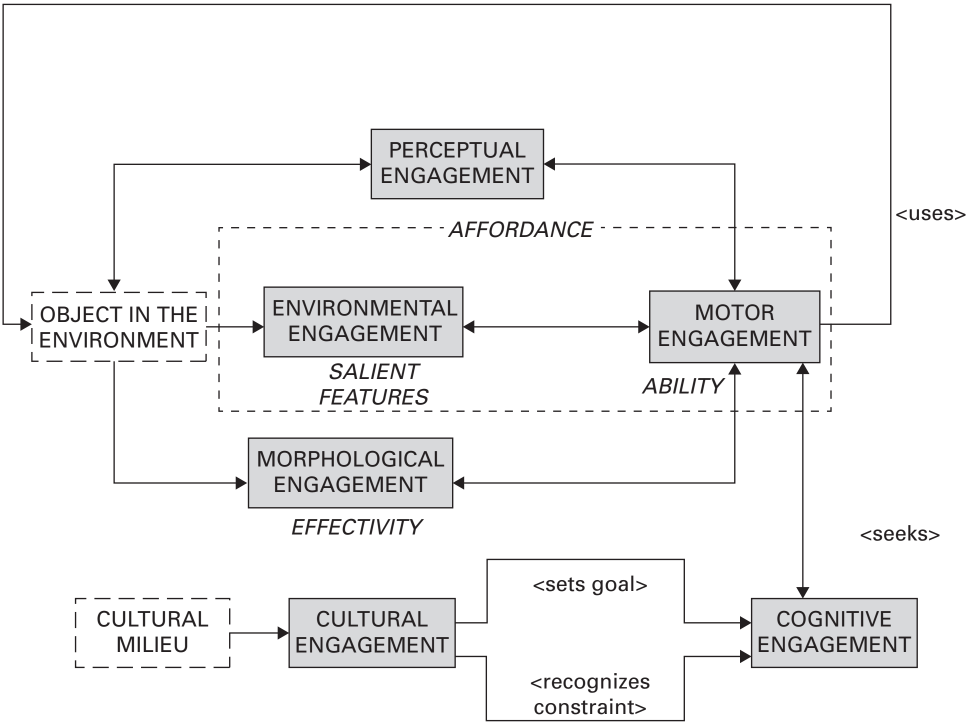

Figure 8.3
Forms of engagement.

Figure 8.3 separates the effectivity of the person, in terms of morphological engagement, which influences how an artifact is grasped, from the ability to reach for or manipulate that artifact, in terms of motor engagement. There are two reasons for this: first, morphology is partly dispositional, for example, in terms of the size of the hand; and second, hand shaping will be influenced by subsequent actions—for example, when reaching to grasp an artifact, hand shape is modified in anticipation of the type of grip required to respond to properties of the artifact, such as weight, fullness, slipperiness, and the like, and this will also be influenced by motor engagement, seeking, for example, "end-state comfort."[13] The assumption is that there is a set of ways in which an artifact can be grasped by the human hand and that the selection of grasp combines artifact properties with intended action. That is, a hand of a given size will have limits as to how it can grasp artifacts, but how the grasp is performed reflects the ability and intentions of the person, which will vary according to a host of situational factors, as well as prior experience.

In order to act on an artifact, there is a need to respond to the "information" that it conveys; that is, information-as-context constrains environmental engagement. Consequently, an affordance arises as the result of the relationship between features perceived through environmental engagement and action performed using motor engagement. We can directly relate this proposal to the formal descriptions of affordance—for instance, imagine we are interested in stair-climbing, and the property of the world is the height of a stair riser, and the property of the person is their leg length. This "perception-action coupling" is the specific relationship between artifact and action and is an emerging property of the world-artifact-person system. However, the relationship is bounded by the other forms of engagement. The suggestion that motor engagement is directed toward subsequent action implies an intention, but I argue that there is equal scope that the "intention" can be defined in response to the motor engagement (i.e., as opportunistic or situated action). At the very least, there is a two-way exchange between the action-as-performed and the goal-state of that action. The role of cognitive engagement is to provide this high-level management to ongoing actions. Across the various forms of engagement, perceptual engagement relates salient features to changing states of the artifact-person system. Finally, the notion of an "acceptable" goal could relate to the culture in which one is acting. This cultural engagement relates to the idea of "cultural affordances"[14] and, as Davis[15] has so clearly explained, carries with it a host of political and social markers. Forms of engagement is still being developed and needs, for example, consideration of “effective engagement” (which was originally assumed to be part of “cultural engagement” but needs it own description).

Activity and Intent

When should a person stop an activity to “off-load” tasks, or when should the agent interrupt a person’s activity with the results of its search? Neither of these issues is trivial, because both depend on the context in which the activity is being performed. From this perspective, the use of technology to support embodied cognition could require the explicit ability (in the technology) to recognize what a person is doing, when interruptions are permissible, how much information to provide, and so on, all of which places the locus of control in the technology.

When we design technology to manage information-as-content, the illusion of “locus of control” is not difficult to maintain. The human issues a request, and the computer provides a response, even if, hidden in this interaction, there is “smart” technology that interprets the request or that determines a request that matches a profile of the requester (as with recommender systems). In this case, it might *feel* as if “agency” lies with the human, but this might simply be because we assume a cause-effect relationship between request and response and see none of the mediation that occurs. In the examples in which the “smart” technology initiates an action because it anticipates or predicts what the user requires, do we assume that it agency is ceded to it? Certainly, we might complain that these smart technologies are operating outside of our control. This might be because there is a lack of transparency in what data they are using or what algorithms they are running; in some cases, we might not even be aware that we are being subjected to data collection. Equally, the smart technology might perform actions that, while they accord with its algorithms, make little sense to the people affected by them. Under these circumstances, speaking of human-computer *interaction* makes little sense because there is no opportunity for “interaction.” In these situations, our relationship with technology has the hallmark of the irony of automation: we are the object of the computer’s algorithm and the subject of its data harvesting but have little ability to manage either. However, the consequences of an action that we perform that has been indicated by the computer may well be our responsibility (rather than the responsibility of either the computer or its manufacturer).

If we accept that “intention” arises from the state of the system, then it makes much more sense to consider ways in which to communicate this intention (as opposed to the immediate inputs and outputs of information-processing). In other words, it might be more useful to know under what constraints information is being collected and processed—for example, what are the goals that are being pursued and who benefits from these? This approach would shift the focus from moment-by-moment data processing and toward the definition of boundaries of the system (who and what are affected?, what are the upper and lower limits of the system’s activity?) and borders (where and how will information be exchanged?). In our everyday life, our interactions with other people are bounded and bordered by a host of conventions that we take for granted (and become aware of only when we encounter situations that are unfamiliar to us). We modify our behavior in response to cues from the situation, in terms of how people are behaving and speaking, for example, or how our conversation partner is responding to what we are saying, and so on. None of these cues is marked explicitly, so they do not constitute information-as-content; rather what they do is help to shape the ecological niches in which our activity occurs. In much the same way, the design of, say, supermarket aisles and product displays is intended to provide information-as-context to encourage ways of moving through the store and “impulse buying” as much as information-as-content in labels, price tags, and signage. We probably pay very little attention to the information-as-context, and few of us challenge (or know how to challenge) the layout of the supermarket because our interaction with it is implicit. However, in the supermarket, we still have the opportunity to act in a way that allows us to retain the sense of agency; we continue to believe that we choose which aisle to walk down and which products to buy. At present, smart technology can often feel as though it has taken Negroponte’s concept of the smart refrigerator to a *reductio ad absurdum*, as if the fridge not only decides what to buy, but makes the order and then expects us to pay for it and put it away when the delivery arrives.

What Is Wrong with “User-Centered” Design?

Design thinking shares with HCI, ergonomics, and systems engineering a recognition that design ought to be “user-centered.” Indeed, to suggest that design might be anything but centered on the user might feel like heresy. It is not my intention to claim that the "user" should not be considered, or have a role, in the practice of design, or that products should not be "usable." But I do question whether "user-centered" simply means asking the user what they want. Such concerns have been voiced by many researchers in the past.[16] There are many well-known and oft-repeated problems that arise from relying solely on user opinion to define requirements for design. For example, users might base their requirements on what they know, or what other products they have seen (e.g., Henry Ford's observation that people wanted a faster horse rather than the automobile). Users struggle to provide clear, unambiguous, and consistent definitions of their actual need; so, they change their minds or contradict themselves, or add to their list of requirements (particularly when they are shown the initial design concept). Add to these the challenge of deciding "who" the user of a product might be, and you appreciate why "user-centered" might be a fig leaf to cover either formal methods (in systems engineering) or the autocratic author (in product design). But user-centered design (and the need to engage with users) is a mandatory part of HCI, with International Standards prescribing a need for this. For example, ISO 13407 defines "usability" partly in terms of measures (efficiency, effectiveness, and user attitude when using a product to perform a task) and partly in terms of "context of use." The latter involves the combination of user (in terms of knowledge, skills, abilities), goal to be achieved, tasks required to achieve the goal, environmental conditions, and other equipment to be used. This notion of context of use chimes with many aspects of design thinking, and the majority of the methods that are advocated in design thinking, or HCI, or ergonomics tends to be more concerned with capturing these elements of a context of use than they are with creativity or generating design concepts.

The idea of context of use also leads to Bardram's notion of activity-based computing. From this, the essential aspects of defining user requirements lies in observing the activity that people do in the environment in which they will be using the product. How one conducts such observations is the basis of the different methods and part of the turf war between disciplines. For ergonomists, the primary approaches involve either task analysis (with its roots in time-and-motion study and information-processing psychology) or cognitive work analysis (chapter 5). A criticism leveled against these approaches is that they are overly reductionistic and fail to capture the subtleties of the context.

Against these "formal" methods, ethnographic approaches produce rich descriptions of the context in which people experience their technology (this approach is considered in chapter 3). We have noted that ethnomethodology in HCI has sought to produce rich descriptions of this complexity, but we have also argued that the reliance on verbal descriptions could miss some of the ways in which the human-artifact-environment system operates. This omission is, as argued above, due to the limits of using words to describe the system. It is also due to the challenge of taking specific instances that have been observed and working from these to general principles. Admittedly, ethnomethodologists might be reluctant to accept that their work *could* be generalized; after all, the richness of the description arises from the specific instance. But, placing this work in the domain of HCI (or design more generally) means that there must be an intention for designers and developers to respond to this work, and this (irrespective of the analysts' intention) means that the results of such studies would be translated into requirements, specifications, and design objectives. In other words, we need to have a "guide for discovery."

Both research traditions focus on capturing the context, but neither offers substantial support for the practice of design. Indeed, both approaches produce documents (complex diagrams from ergonomics, reportage from ethnography), which are more or less handed over to the designer to interpret. HCI has been divided between approaches to evaluation that emphasize "usability" (covered by the International Standards and with its root it ergonomics) and "user experience" (with roots that can be traced to the ethnographic approaches used in HCI). As noted above, ISO 13407 would expect both approaches (user attitude is, after all, a reflection of user experience).

To return to our discussion, the question is what does it mean to "center" design on users? If asking them what they want is problematic or observing them in situ is either reductionistic or superficial, what is left? My impression is that the experienced designer seeks to define a particular situation (or "context of use"). The description of the context of use can be obtained and reviewed by speaking to people, watching them, experiencing the situation oneself, or creating scenarios, stories, or prototypes. Each of these broad approaches (and the countless methods associated with them) is used to recast the context of use into terms that are amenable to design. The methods that are used (and the baggage that they bring with them) become less important than the overarching design to reach consensus on this context of use. What shifts design from a mundane task of problem-solving to something more difficult is that the problems are typically "wicked" and their solution involves optimizing multiple objectives (not least because the idea of a "consensus" can be hugely challenging). By understanding context of use as the ecological niche for the potential users of the artifact and by appreciating how the affordances offered by the artifact create opportunities for action, design practice involves the broad and deep appreciation of the potential states in which specific human-artifact-environment systems can function. In order to do understand context of use, accounts that draw on phenomenology and ethnography (e.g., through reportage, conversation analysis, video analysis) provide the breadth of understanding. However, in order to achieve the depth, it is also important to remember Brunswik's notion of "ecological validity" and to specify the ways in which the environment and activity acquire salience. My proposal is that this requires analysis that can capture the micro-materialities of states of the human-artifact-environment system (e.g., through data from eye-tracking or sensors on the person or the artifact). This fine-grained analysis of micro-materialities is not intended to reflect discrete actions (which would, I think, do little to complement the detail in phenomenological accounts) but should be analyzed in terms of the stability of the system—that is, in terms of how the balance between variability and consistency is maintained by the application of metrics that describe entropy and that use techniques from dynamic systems. This marriage of phenomenology and dynamic systems is at the very heart of RECS and, I propose, offers opportunities for development of digital technology (in terms of recognizing activity and intent as well as in terms of more deeply understanding human interaction with these technologies).

HCI has been concerned with the lived experience of everyday interactions with digital technology, or what Dourish calls "everyday mundane experience."[17] Dourish's approach derives from particular readings of phenomenology (drawing principally on aspects of the work of Heidegger and Merleau-Ponty). For Dourish,[18] in his pioneering work on embodiment in human-computer interaction, meaning involves three elements. The first two develop from his reading of Merleau-Ponty and his insistence on "intentionality" and "ontological commitment." In terms of the former, Dourish uses the phrase "coupling" to elucidate this: "By coupling, I mean the way that we can build up and break down relationships between entities, putting them together or taking them apart for the purpose of incorporating them into our action."[19] I understand this idea of coupling as being similar to Gibson's "complementarity" or Ingold's "co-respondance." which relate to the experience of purposefully interacting within an ecological niche. How we recognize and realize the material and functional environment can depend on the constraints within which we act. So, there would be differences between the experiences of a designer and user of a product; both can approach the same physical form from different ontological commitments. At issue, is the extent to which the material or the functional environment becomes influenced by the form or the function of the artifact, and how this influences action.

What Can Design and Creativity Tell Radical Embodied Cognitive Science?

For its critics "representation-hungry" cognition, such as abstract thinking and creativity, lie outside the realms of embodied cognition. The argument is that, while embodied approaches can deal with physical actions (and while there might a grudging admission that embodied cognition can provide reasonable accounts of problem-solving tasks that rely on physical artifacts), it is unable to cope with "higher-order" cognition. In part this argument relies on the assumption that there needs to be internal representation in order to "do" cognition. These "higher-order" cognitive activities, for information processing, require the creation and manipulation of internal representations that are built on top of existing internal representations. The tautology of the argument (that internal representations are needed to create internal representations) undermines the argument, but still, if there is no need for internal representation, then how can this higher-order cognition occur?

I believe that RECS provides a clear description of creative practice that feels as if it has far more in common with the experience of practitioners than the information-processing approaches. Creativity requires a repertoire of responses (in terms of the ability to work with materials, to produce models, to make sketches, and so on) that are acquired from the practice and experience of designing. These responses provide a means of interpreting and responding to constraints within the problem space. My reading of

RECS sees creativity as the opportunistic response to constraints. Indeed, creativity seems to thrive with constraints; that is, people are more likely to produce solutions that are rated as "creative" when there are *some* constraints than when there are no constraints.[20]

In a similar vein, for the skilled intentionality framework (SIF), higher-order cognition does not depend on internal representations. In part this is because human activity is performed in a "rich landscape of affordances," which are responded to on the basis of the lived experience of each of us. Recall that affordances, in this book, relate to states of the human-artifact-environment system that can be considered as objectives for transitions to and from. First, the environment (as a socio-material construct) has a host of constraints on what actions are performable (in physical terms) and permissible (in outcome terms). Second, the notions of performable and permissible will be informed by the lived experience of each individual. If we share similar cultural upbringing, then we might also share similar notions of what is permissible. The notion of creativity advanced in this book is one in which activity seeks to discover the limits of the constraints (through tightening or loosening one or two of these to be objectives for the system's transitions). In the words of Withagen and van der Kamp, "Creativity can be conceived as the discovery and creation of unconventional affordances."[21] This could easily become a trial-and-error form of activity (or, in artificial intelligence terms, a form of "explore and exploit" activity). It should not, I think, be seen as aimless, reckless, or unstructured—because the precursor to such activity is the identification of constraints to modify. So, in part, higher-order cognition arises from identifying and manipulating the constraints in the human-artifact-environment system. As new states of the system arise, so these need to be evaluated (in terms of affordances or in terms of permissible outcomes). None of this requires internal representation. But advocates of information-processing approaches might feel that this account is overly mechanistic and that it does not provide an indication of how things such as "insight" might arise.

Notwithstanding that many advocates of the information-processing approach dislike the notion of "insight" (because it does not conform to the production line of cognition that the approach assumes), I feel that this complaint misses two essential issues. First, RECS (and SIF and other approaches to embodied cognition) recognizes that human activity is part of the ongoing, lived experience of the individual. For RECS, an activity not only occurs in the human-artifact-environment system but does so in terms of the history of the individual's history of prior interactions. This is partly in terms of solutions to the degrees of freedom problem that has previously been applied, both in terms defining of physical movement in terms of coordinative structures and in terms of selecting and responding to sense data, e.g., in terms of information-as-context. It is also partly in terms of the history of the environment and the artifact (which might be modified as the result of previous activities upon them). This activity is not and cannot be (as information-processing approaches so often imply) discrete, boxed off, and separated from our history of activity. In the laboratory, are we really to assume that the participant forgets everything she knows and focuses only on the task instructions and the task materials in front of her? Even if this were the case, are we sure that these task materials provide sufficient ecological validity to be a fair test of the whatever activity under investigation? In terms of "insight" (as we saw in chapter 2) a reasonable explanation (and one that does not require us to conjure solutions to a problem out of thin air) is that participants respond to affordances in the context of the problem-solving task and that these create opportunities for action. Similarly, epistemic action in problem-solving capitalizes on this activity. Whether the environment is the design studio, the jeweler's workshop, the keyboard of a piano, or the canvas on an easel, the activity of the human responds to its properties and changes it. In this way, creativity is simply the dynamic interplay between the artist or designer and their environment. Or rather, it is the continual reciprocal engagement between elements in the human-artifact-environment system.

Notes

Preface

1. Baber, C., 1991, *Speech Technology in Control Room Systems: A Human Factors Perspective*, Chichester: Ellis Horwood.

2. Baber, C., and Stanton, N., 2017, Rewritable routines in human interaction with public technology, *Engineering Psychology and Cognitive Ergonomics: Volume 2: Job Design and Product Design*, Avebury: Ashgate.

3. Baber, C., and Butler, M., 2012, Expertise in crime scene examination: Comparing search strategies of expert and novice crime scene examiners in simulated crime scenes, *Human Factors*, *54*, 413–424.

4. Johannessen, C. M., Longcamp, M., Stuart, S. A., Thibault, P. J. and Baber, C., 2021, The look of writing in reading. Graphetic empathy in making and perceiving graphic traces. *Language Sciences*, *84*.

5. Baber, C., 2003, *Cognition and Tool Use*, London: Taylor and Francis.

6. http://avant.edu.pl/en/events-2.

7. Baber, C., and Janulis, K., 2021, Purposeful tool use in early lithic technologies, *Adaptive Behavior*, *29*, 169–180.

8. Chen, X., Starke, S. D., Baber, C. and Howes, A., 2017, A cognitive model of how people make decisions through interaction with visual displays, *Proceedings of the 2017 CHI conference on human factors in computing systems*, 1205–1216, New York: Association of Computing Machinery.

9. Baber, C., Chemero, T. and Hall, J., 2019, What the jeweller's hand tells the jeweller's brain: Tool use, creativity and embodied cognition, *Philosophy & Technology*, *32*, 283–302.

Chapter 1

1. Putnam, H., 1975, The meaning of “meaning,” *Philosophical Papers, Vol. II: Mind, Language, and Reality*, Cambridge: Cambridge University Press.

2. Shipp, N., and Vallee-Tourangeau, F., 2019, Cognition can be distributed, extended, enacted, embodied and systemic (but does it matter which?), *British Psychological Society Cognitive Section Bulletin 4*.

3. Shapiro, L., 2011, *Embodied Cognition*, New York: Routledge.

4. Wilson, M., 2002, Six views of embodied cognition, *Psychonomic Bulletin and Review 9*, 625–636.

5. James, W., 1890, *The Principles of Psychology*, New York: Holt, 1:154.

6. Vygotsky, L. S., 1962, *Thought and Language*, Cambridge, MA: MIT Press.

7. Wertsch, J. V., 1988, *Vygotsky and the Social Formation of Mind*, Cambridge, MA: Harvard University Press.

8. Brooks, R., 1991, Intelligence without representation, *Artificial Intelligence Journal*, *47*, 139–160; Chemero, A., 2009, *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press; Dreyfus, H. L., 2002, Intelligence without representation–Merleau-Ponty’s critique of mental representation: The relevance of phenomenology to scientific explanation, *Phenomenology and the Cognitive Sciences*, *1*, 367–383; Gallagher, S., 2017, *Enactivist Interventions: Rethinking the Mind*, Oxford: Oxford University Press; Hutto, D. D., and Myin, E., 2013, *Radicalizing Enactivism: Basic Minds without Content*, Cambridge, MA: MIT Press.

9. Sennett, R., 2008, *The Craftsman*, London: Penguin, 227.

10. Bateson, G., 1971, *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology*, Chicago: University of Chicago Press, 5.

11. Williams, D., 2018, Predictive processing and the representation wars, *Minds and Machines, 28, 141–172*.

12. Dietrich, E., 2007, Representation, in P. Thaggard, ed., *Philosophy of Psychology and Cognitive Science*, Oxford: North-Holland, 1–29.

13. Baddeley, A., 2007, *Working Memory, Thought, and Action*, Oxford: Oxford University Press.

14. Miller, G. A., 1956, The magical number seven, plus or minus two: Some limits on our capacity for processing information, *Psychological Review*, *63*, 81.

15. Baddeley, A. D., Thomson, N., and Buchanan, M., 1975, Word length and the structure of short-term memory, *Journal of Verbal Learning and Verbal Behavior*, *14*, 575–589.

16. Fodor, J. A., and Pylyshyn, Z. W., 1988, Connectionism and cognitive architecture: A critical analysis, *Cognition* 28, 3–71.

17. O'Brien, G., and Opie, J., 2004, Notes toward a structuralist theory of mental representation, in H. Clapin, P. Staines and P. Slezak, eds., *Representation in Mind: New Approaches to Mental Representation*, 1–20, Amsterdam: Elsevier.

18. Markman, A. B., and Dietrich, E., 2000, In defense of representation, *Cognitive Psychology*, *40*, 140.

19. Craik, K. J. W., 1952, *The Nature of Explanation*, Cambridge: Cambridge University Press, 61.

20. Craik, K. J., 1947, Theory of the human operator in control systems 1: I. The operator as an engineering system, *British Journal of Psychology. General Section*, *38*, 56–61; Craik, K. J., 1948, Theory of the human operator in control systems. II. Man as an element in a control system, *British Journal of Psychology*, *38*, 142.

21. Bartlett, F. C., 1932, *Remembering: A Study in Experimental and Social Psychology*, Cambridge: Cambridge University Press.

22. Craik, 1952, 51.

23. Walter, W. G., 1953, *The Living Brain*, London: Butterworth, 125.

24. Gibson, J. J., 1966, *The Senses Considered as Perceptual Systems*, Boston: Houghton Mifflin.

25. Merleau-Ponty, M., 2014 [1945], *Phenomenology of Perception*, London: Routledge.

26. Chemero, A., 2020. Epilogue: What Embodiment Is. In N.K. Dees (ed) A Multidisciplinary Approach to Embodiment, London: Routledge, 133–139.

27. Rorty, 1979. *Philosophy and the Mirror of Nature*, Princeton, NJ: Princeton University Press.

28. Gibbs, R. W., 2006, *Embodiment and Cognitive Science*, Cambridge: Cambridge University Pres; Rupert, R. D., 2009, *Cognitive Systems and the Extended Mind*, Oxford: Oxford University Press.

29. Gallagher, 2017, 160.

30. Dreyfus, H., 1991, *What Computers Still Can't Do*, Cambridge, MA: MIT Press.

31. Polyani, M., 1966, *The Tacit Dimension*, London: Routledge and Kegan Paul.

32. Hutto and Myin, 2013, 67.

33. Rouse, W. B., and Morris, N. M., 1986, On looking into the black box: Prospects and limits in the search for mental models, *Psychological Bulletin*, *100*, 349.

34. Ryle, G., 1949, *The Concept of Mind*, London: Hutchinson.

35. Gallagher, S., 2015, How embodied cognition is being disembodied, *Philosophers" Magazine*, *68*, 96–102.

36. Butler, S., 1912, *The Notebooks of Samuel Butler*, London: A. C. Fifield.

37. Varela, F. J., 1979, *Principles of Biological Autonomy*, New York: North Holland.

38. Chomsky, N., 1988, *Language and the Problems of Knowledge*, Cambridge, MA: MIT Press.

39. Shannon, C. E., 1948, A mathematical theory of communication, *Bell System Technical Journal*, *27*, 379–423, 623–656.

40. Turvey, M. T., and Carello, C., 1985, The equation of information and meaning from the perspectives of situation semantics and Gibson's ecological realism, *Linguistics and Philosophy*, *8*, 81–90.

41. Gibson, J. J., 1979, *The Ecological Approach to Visual Perception*, Boston, MA: Houghton Mifflin, 242.

42. Merleau-Ponty, 2014.

43. Aleksander, I., 2013, *The World in My Mind, My Mind in the World*, London: Andrews UK Limited

44. Kirsh, D., 2010, Thinking with external representations, *AI & Society*, *25*, 441–454; Scaife, M., and Rogers, Y., 1996, External cognition: How do graphical representations work? *International Journal of Human-Computer Studies*, *45*, 185–213.

45. Hollan, J., Hutchins, E., and Kirsch, D., 2002, Distributed cognition: toward a new foundation for human-computer interaction, in J. Carroll, ed., *Human-Computer Interaction in the New Millennium*, 75–94, New York: Addison-Wesley; Hutchins, E., 1995, *Cognition in the Wild*, Cambridge, MA: MIT Press.

46. Thelen, E., and Smith, L. B., 1994, *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge, MA: MIT Press; Van Gelder, T., 1995, What might cognition be, if not computation? *The Journal of Philosophy*, *92*, 345–381.

47. Hutchins, 1995.

48. Nemeth, C., 2003, How cognitive artifacts support distributed cognition in acute care, *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society*, 381–385, Santa Monica, CA: Human Factors And Ergonomics Society.

49. Clark, A., and Chalmers, D. J., 1998, The extended mind, *Analysis*, *58*, 10–23.

50. Adams, F., and Aizawa, K., 2008, *The Bounds of Cognition*, Oxford: Blackwell.

51. Clark, A., 2001, Reasons, robots, and the extended mind, *Mind and Language*, *16*, 132.

52. Brooks, R., 1990, Elephants don"t play chess, *Robotics and Autonomous Systems*, *6*, 3–15.

53. Goldinger, S. D., Papesh, M. H., Barnhart, A. S., Hansen, W. A., and Hout, M. C., 2016, The poverty of embodied cognition, *Psychonomic Bulletin and Review, 23*, 959–978.

54. Clark, A., and Toribio, J., 1994, Doing without representing? *Synthese, 101*, 418.

55. Degenaar, J., and Myin, E., 2014, Representation—hunger reconsidered, *Synthese*, *191*, 3639–3648.

56. Barsalou, L. W., 1999, Perceptions of perceptual symbols, *Behavioral and Brain Sciences*, *22*, 637–660; Barsalou, L. W., 2008, Grounded cognition, *Annual Review of Psychology*, *59*, 617–645.

57. Calvo-Merino, B., Glaser, D. E., Grèzes, J., Passingham, R. E., and Haggard, P., 2005, Action observation and acquired motor skills: An FMRI study with expert dancers, *Cerebral Cortex*, *15*, 1243–1249.

58. Dreyfus, 2002, 382.

59. Gallagher, 2017.

60. Bernstein, N. A., 1967, *The Coordination and Regulation of Movements*, Oxford: Pergamon Press.

61. Gibson, 1979, 225.

62. Huys et al., 2004, 360.

63. Luria, A. R., 1973, *The Working Brain*, Harmondsworth: Penguin.

64. Keele, S. W., 1968, Movement control in skilled motor performance, *Psychological Bulletin*, *70*, 387; Schmidt, R. A., 1975, A schema theory of discrete motor skill learning, *Psychological Review*, *82*, 225.

65. Gallagher, S., 2005, *How the Body Shapes the Mind*, Oxford: Clarendon Press.

66. Wilson, A. D., and Golonka, S., 2013, Embodied cognition is not what you think it is, *Frontiers in Psychology*, *4*, 58.

67. Barsalou, 1999, 2008; Lakoff, G., and Johnson, M., 1980, *Metaphors We Live By*, Chicago: University of Chicago Press.

68. Lakoff and Johnson, 1980.

69. Chemero, 2009, 29.

70. Gallagher, S., and Zahavi, D., 2008, *The Phenomenological Mind: An Introduction to Philosophy of Mind and Cognitive Science*, New York: Routledge.

71. Adams and Aizawa, 2008.

Chapter 2

1. Glaveanu, V. P., 2014, *Distributed Creativity: Thinking outside the Box of the Creative Individual*, Cham, Germany: Springer; Malinin, L. H., 2019, How radical is embodied creativity? *Frontiers in Psychology*, *10*, 2372.

2. Malinin, 2019.

3. "Kees Overbeeke," Interaction Design Foundation, n.d., https://www.interaction-design.org/literature/author/kees-overbeeke.

4. Rowe, 1991, *Design Thinking*, Cambridge, MA: MIT Press, 2.

5. Rowe, 1991, 77.

6. Csikszentmihalyi, M., 1996, *Flow and the Psychology of Discovery and Invention*, New York: Harper Perennial.

7. Dreyfus, H. L., 2002, Intelligence without representation–Merleau-Ponty's critique of mental representation: The relevance of phenomenology to scientific explanation, *Phenomenology and the Cognitive Sciences*, *1*, 367–383.

8. Gordon, W. J. J., 1961, *Synectica: The Development of Creative Capacity*, New York: Harper and Row.

9. Koestler, A. 1964, *The Act of Creation*, New York: Macmillan.

10. Zeng, L., Proctor, R. W., and Salvendy, G., 2011, Can traditional divergent thinking tests be trusted in measuring and predicting real-world creativity?, *Creativity Research Journal, 23*, 24–37.

11. Gallagher, S. 2005, *How the Body Shapes the Mind*. Oxford: Clarendon Press.

12. Malinin, 2019; Sawyer, R. K., and de Zutter, S., 2009, Distributed creativity: How collective creations emerge from collaboration, *Psychology of Aesthetics, Creativity, and the Arts*, *3*, 81; Schön, D. A., 1983, *The Reflective Practitioner: How Professionals Think in Action*, London: Routledge; Csikszentmihályi, M., 1990, The domain of creativity, in M. A. Runco and R. S. Albert, eds., *Sage Focus Editions, Vol. 115: Theories of Creativity*, 190–212, London: Sage.

13. Boden, M. A., 1996, What is creativity? In M. A. Boden, ed., *Dimensions of Creativity*, Cambridge, MA: MIT Press, 75.

14. Brown, T., 2009, *Change by Design: How Design Thinking Transforms Organizations and Inspires Innovation*, New York: Harper Business.

15 .Simon, H. A., 1988, The science of design: creating the artificial, *Design Issues*, *4*, 67–82.

16. Asimow, M., 1962, *Introduction to Design*, Englewood Cliffs, NJ: Prentice-Hall.

17. Kimbell, L., 2011, Rethinking design thinking: Part I. *Design and Culture*, *3*, 285–306; Kimbell, L., 2012, Rethinking design thinking: Part II. *Design and Culture*, *4*, 129–148.

18. Campbell, D. T., 1960, Blind variation and selective retention in creative thought as in other knowledge processes, *Psychololgical Review*, *67*, 380–400.

19. Voss, J. F., 2006, Toulmin's model and the solving of ill-structured problems, in D. Hitchcock and B. Verheij, eds., *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation*, 300–311, Berlin: Springer.

20. Simon, H. A., 1973, The structure of ill-structured problems, *Artificial Intelligence*, *4*, 181–201.

21. Simon, H. A., 1972, Theories of bounded rationality, *Decision and organization*, *1*, 161–176.

22. Ashton, K., 2015, *How to Fly a Horse: The Secret History of Creation, Invention, and Discovery*, London: Doubleday, 59.

23. Reitman, W. R., 1964, Heuristic decision procedures, open constraints and the structure of ill-defined problems, in M. Shelly II and G. L. Bryan, eds., *Human Judgments and Optimality*, 282–315, New York: John Wiley & Sons.

24. Stokes, P. D., 2006, *Creativity from Constraints: The Psychology of Breakthrough*, New York: Springer

25. Reitman, 1964, 307.

26. Baber, C., Chemero, T., and Hall, J., 2019, What the jeweller's hand tells the jeweller's brain: Tool use, creativity and embodied cognition, *Philosophy and Technology, 32*, 283–302.

27. Kirsh, D., and Maglio, P., 1994, On distinguishing epistemic from pragmatic action, *Cognitive Science*, *18*, 513–549, quoted in the caption to figure 2.1.

28. Kirsh and Maglio, 1994.

29. Maier, N. R. F., 1931, Reasoning in humans II: The solution of a problem and its appearance in consciousness, *Journal of Comparative and Physiological Psychology, 12*, 181–194.

30. Maier, N. R. F., 1931, Reasoning in humans II: the solution of a problem and its appearance in consciousness, *Journal of Comparative and Physiological Psychology, 12*, 193.

31. Ohlsson, S., 1992, Information-processing explanations of insight and related phenomena, *Advances in the Psychology of Thinking, 1*, 1–44.

32. Cofer, C. N., 1951, Verbal behaviour in relation to reasoning and values, in H. Guetzkow, ed., *Groups, Leadership and Men*, 206–217, Pittsburgh, PA: Carnegie Press.

33. Duncan, C. P., 1961, Attempts to influence performance on an insight problem, *Psychological Reports*, *9*, 35–42.

34. Olteteanu, A-M., and Freksa, C., 2014, Towards affordance-based solving of object insight problems, *Proceedings of the 1st Workshop on Affordances: Affordances*

for Vision in for Cognitive Robotics, Robotics Science and Systems, Carnegie Music Hall, Pittsburgh, PA: Robotics: Science and Systems.

35. Zhang, J., and Norman, D., 1994, Representations in distributed cognitive tasks, *Cognitive Science, 18*, 87–122.

36. Kotovsky, K., Hayes, J. R., and Simon, H. A., 1985, Why are some problems hard? Evidence from the Tower of Hanoi, *Cognitive Psychology*, *17*, 248–294.

37. Guthrie, L. G., Vallée-Tourangeau, F., Vallée-Tourangeau, G., and Howard, C., 2015, Learning and interactivity in solving a transformation problem, *Memory & Cognition*, *43*, 723–735.

38. Vallée-Tourangeau, F., Steffensen, S. V., Vallée-Tourangeau, G., and Sirota, M., 2016, Insight with hands and things, *Acta Psychologica*, *170*, 195–205.

39. Steffensen, S. V., Vallée-Tourangeau, F. and Vallée-Tourangeau, G., 2016, Cognitive events in a problem-solving task: a qualitative method for investigating interactivity in the 17 Animals problem. *Journal of Cognitive Psychology*, *28*, 79–105.

40. Archer, L. B., 1965, *Systematic Method for Designers*, London: Her Majesty's Stationery Office.

41. Jones, J. C., 1992, *Design Methods*, 2nd ed., New York: van Nostrand Reinhold, 19.

42. Polyani, M., 1966, *The Tacit Dimension*, London: Routledge and Kegan Paul.

43. Perry, S., 2018, *Making Sense: Cognition, Computing, Art and Embodiment*, Cambridge, MA: MIT Press, 254.

44. Adamson, G., 2013, *Thinking through Craft*, London: Bloomsbury, 4.

45. Baber et al., 2019.

46. Rietveld, E., and Brouwers, A. A., 2016, Optimal grip on affordances in architectural design practices: An ethnography, *Phenomenology and Cognitive Sciences*, *16*, 555.

47. Verstijnen, I. M., van Leeuwen, C., Goldschmidt, G., Hamel, R., and Hennessey, J. M., 1998, Creative discovery in imagery and perception: Combining is relatively easy, restructuring takes a sketch, *Acta Psychologica*, *99*, 177–200.

48. Goldschmidt, G., 1991, The dialectics of sketching, *Creativity Research Journal*, *3,* 123–143.

49. Alexander, C., 1964, *Notes on the Synthesis of Form*, Cambridge MA: Harvard University Press.

50. Baber et al., 2019.

51. Chemero, A., 2009, *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press

52. Rietveld and Brouwers, 2016.

53. Baber, C. and Saini, M., 1995, Craft skills in jewellery manufacture, in S. A. Robertson (ed), *Contemporary Ergonomics 1995*, 92–97, London: Taylor and Francis.

54. Pereira, Á., and Tschimmel, K., 2012, The design of narrative jewelry as a perception in action process, in Duffy, A.; Nagai, Y.: Taura, T., eds, *DS 73–1 Proceedings of the 2nd International Conference on Design Creativity Volume 1.*, Glasgow: The Design Society.

55. Baber et al., 2019.

56. Ingold, T., 2010, The textility of making, *Cambridge Journal of Economics, 34*, 91–10; Ingold, T., 2013, *Making: Anthropology, Archaeology, Art and Architecture*. New York: Routledge.

57. Schön, 1983.

58. Schön, 1983, 50.

59. Dewey, J., 1933, *How We Think*, Boston: Houghton Mifflin.

60. Schön, 1983, 103.

61. Wuytens, K., and Willems, B., 2009, Diversity in the design processes of studio jewellers, *EKSIG: Experimental Knowledge, Method & Methodology*, http://www.academia.edu/885664/Diversity_in_the_design_processes_of_studio_jewellers.

Chapter 3

1. Gibson, J. J., 1979, *The Ecological Approach to Visual Perception*, Boston: Houghton Mifflin, 128.

2. Gibson, 1979, 128.

3. Rietveld, E., and Kiverstein, J., 2014, A rich landscape of affordances, *Ecological Psychology*, *26*, 325–352.

4. Ingold, T., 2000, *The Perception of the Environment: Essays on Livelihood, Dwelling and Skill*, London: Psychology Press, 5.

5. Gibson, 1979, 254.

6. Freivalds, A., 1986, The ergonomics of shovelling and shovel design—A review of the literature, *Ergonomics*, *29*, 3–18.

7. Rietveld, E., 2008, Situated normativity: The normative aspect of embodied cognition in unreflective action, *Mind*, *117*, 973–1001.

8. Newell, K. M., 1986, Constraints on the development of coordination, in M. G. Wade and H. T. A. Whiting, eds., *Motor Development in Children: Aspects of Coordination and Control*, 341–361, Amsterdam: Martinus Nijhoff Publishers.

9. Gibson J. J., 1947, Motion picture testing and research. *AAF Aviation Psychology Research Report 7*. Washington, DC: US Government Printing Office.

10. Garcia-Marquez, G., 1967, *One Hundred Years of Solitude*, London: Picador, 46.

11. Gibson, J. J., and Crooks, L. E., 1938, A theoretical field-analysis of automobile-driving, *American Journal of Psychology*, *51*, 453–471.

12. Flach, J. M., 1995, The ecology of human machine systems: A personal history, in J. M. Flach, P. A. Hancock, J. Caird, and K. J. Vicente, eds., *Global Perspectives on the Ecology of Human-Machine Systems*, 121–156, Boca Raton, FL: CRC Press.

13. Brunswik, E., 1956, *Perception and the Representative Design of Psychological Experiments*, Berkeley: University of California Press.

14. Vicente, K. J., 2003, Beyond the lens model and direct perception: Toward a broader ecological psychology, *Ecological Psychology*, *15*, 241–267.

15. Vicente, 2003.

16. Warren, W. H., Jr., 1984, Perceiving affordances: Visual guidance of stair climbing, *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 683–703.

17. Vicente, 2003.

18. Camerer, C. F., and Johnson, E. J., 1997, The process-performance paradox in expert judgment: How can experts know so much and predict so badly, in A. Ericsson and J. Smith, eds, *Research on Judgment and Decision Making: Currents, Connections, and Controversies*, Cambridge: Cambridge University Press, 195–217.

19. Gigerenzer, G., 2008, Why heuristics work, *Perspectives on Psychological Science*, *3*, 20–29.

20. Karelaia, N., and Hogarth, R., 2008, "Determinants of linear judgment: A meta-analysis of lens model studies," *Psychological Bulletin*, *134*, 404–426.

21. Bisantz, A. M., Kirlik, A., Gay, P., Phipps, D. A., Walker, N., and Fisk, A. D., 2000, Modeling and analysis of a dynamic judgment task using a lens model approach, *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems And Humans*, *30*, 605–616.

22. Kirlik, A., Walker, N., Fisk, A. D., and Nagel, K., 1996, Supporting perception in the service of dynamic decision making, *Human Factors*, *38*, 288–299.

23. Lipshitz, R. and Ben Shaul, O., 1997, Schemata and mental models in recognition-primed decision making, in C. E. Zsambok and G. Klein, eds., *Naturalistic Decision Making*, London: Psychology Press, 293–303.

24. Baber, C., Chen, X., and Howes, A., 2015, (Very) rapid decision-making: Framing or filtering? Communication presented at the International Conference on Naturalistic Decision-Making, McLean, VA.

25. Vallée-Tourangeau, F., Steffensen, S. V., Vallée-Tourangeau, G., and Sirota, M., 2016, Insight with hands and things, *Acta Psychologica*, *170*, 195–205.

26. Klein, G. A., 1993, A recognition-primed decision (RPD) model of rapid decision making, in G. A. Klein, J. Orasanu, R. Calderwood, and C. E. Zsambok, eds., *Decision Making in Action: Models and Methods*, 138–147, Norwood, NJ: Ablex.

27. Weller, A., Villejoubert, G., and Vallée-Tourangeau, F., 2011, Interactive insight problem solving, *Thinking and Reasoning*, *17*, 429–439.

28. Csikszentmihalyi, M., 1996, *Flow and the Psychology of Discovery and Invention*, New York: Harper Perennial

29. Gibson, 1979, 9.

30. Rasmussen, J., 1999, Ecological interface design for reliable human-machine systems, *International Journal of Aviation Psychology*, *9*, 203–223.

31. Suchman, L. A., 1987, *Plans and Situated Action*, Cambridge: Cambridge University Press.

32. Heidegger, M., 1962, *Being and Time*, New York: Harper Row.

33. Clark, P., 1976, Atomism versus thermodynamics, in C. Howson, ed., *Method and Appraisal in the Physical Sciences*, Cambridge: Cambridge University Press, 41–106.

34. Sharrock, W. W., and Anderson, B., 2012, *The Ethnomethodologists*, London: Routledge, 124.

35. Høffding, S. and Martiny, K., 2016, Framing a phenomenological interview: What, why and how, *Phenomenology and the Cognitive Sciences*, *15*, 539–564.

36. Klein, G. A., Calderwood, R., and Macgregor, D., 1989, Critical decision method for eliciting knowledge, *IEEE Transactions on Systems, Man, and Cybernetics*, *19*, 462–472.

37. Gallagher, S., 2017, *Enactivist Interventions: Rethinking the Mind*, Oxford: Oxford University Press.

38. Gallagher, S., and Zahavi, D., 2008, *The Phenomenological Mind: An Introduction to Philosophy of Mind and Cognitive Science*, New York: Routledge, 28.

39. Dourish, P., 2004, *Where the Action Is: The Foundations of Embodied Interaction*, Cambridge, MA: MIT Press, 125.

40. Ericsson, K. A., and Simon, H. A., 1980, Verbal reports as data, *Psychological Review*, *87*, 215.

41. The video *William "Holly" Whyte in His Own Words, "The Social Life of Small Urban Spaces" (1980)* gives a very neat summary of the work; see https://www.youtube.com/watch?v=sU2vVqbtRAY.

42. Rietveld and Kiverstein, 2014.

43. Baber, C., and Butler, M., 2012, Expertise in crime scene examination: Comparing search strategies of expert and novice crime scene examiners in simulated crime scenes, *Human Factors*, *54*, 413–424.

44. Heath, C., and Luff, P., 1992, Collaboration and control: Crisis management and multimedia technology in London Underground Line control rooms, *Journal of Computer-Supported Cooperative Work*, *1*, 24–48; Heath, C., Luff, P., Sanchez-Svensson, M., and Nicholls, M., 2017, Exchanging implements: The micro-materialities of multidisciplinary work in the operating theatre, *Sociology of Health and Illness*, *40*, 297–313.

45. Chen, X., Starke, S. D., Baber, C., and Howes, A., 2017, A cognitive model of how people make decisions through interaction with visual displays, *CHI17: Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems*, New York: ACM, 1205–1216.

46. Hayhoe, M., and Ballard, D., 2005, Eye movements in natural behaviour, *Trends in Cognitive Sciences*, *9*, 188–194.

47. Ballard, D. H., Hayhoe, M. M., and Pelz, J. B., 1995, Memory representations in natural tasks, *Journal of Cognitive Neuroscience*, *7*, 66–80.

48. Moray, N., and Rotenberg I. R. A., 1989, Fault management in process control: Eye movements and action, *Ergonomics*, *32*, 1319–1342; Starke, S. D., Baber, C., Cooke, N. J., and Howes, A., 2017, Workflows and individual differences during visually guided routine tasks in a road traffic management control room, *Applied Ergonomics*, *61*, 79–89.

49. Kirwan, B., Kaarstad, M., Hauland, G., and Follesoe, K., 1995, See no evil, hear no evil, speak no evil: Verbal protocol analysis, eye movement analysis, and nuclear power plant diagnosis, *Contemporary Ergonomics 1995*, London: Taylor and Francis, 249–249.

50. Heath et al., 2017, 300.

51. Heath et al., 2017, 303.

52. Latour, B., and Woolgar, S., 1979, *Laboratory Life: The Social Construction of Scientific Facts*, Beverly Hills, CA: Sage.

53. Dourish, P., 2017, *The Stuff of Bits: An Essay on the Materialities of Information*, Cambridge, MA: MIT Press.

54. Baber, C., Cengiz, T. G., Starke, S. and Parekh, M., 2015, Objective classification of performance in the use of a piercing saw in jewellery making, *Applied Ergonomics*, *51*, 211–221.

Chapter 4

1. Norman, D. A., 1988, *The Psychology of Everyday Things*, New York: Basic Books, 9.

2. Chemero, A., 2020, Epilogue: What embodiment is. In N. K. Dees, ed., *A Multidisciplinary Approach to Embodiment*, London: Routledge, 133–139.

3. Varela, F., Thompson, E., and Rosch, E., 1991 *The Embodied Mind: Cognitive Science and Human Experience*, Cambridge, MA: MIT Press.

4. Baber, C., 2018 Designing smart objects to support affording situations: Exploiting affordance through an understanding of forms of engagement, *Frontiers in Psychology*, *9*, 292.

5. Brooks, R., 1991, Intelligence without representation, *Artificial Intelligence Journal*, *47*, 139–160.

6. Montesano, L., Lopes, M., Bernardino, A., and Santos-Victor, J., 2008, Learning object affordances: From sensory-motor coordination to imitation, *IEEE Transactions on Robotics*, *24*, 15–26.

7. Cisek, P., 2007, Cortical mechanisms of action selection: The affordance competition hypothesis, *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*, 1585–1599.

8. Sridharan, M., and Meadows, B., 2017, An architecture for discovering affordances, causal laws, and executability conditions, *Advances in Cognitive Systems*, *5*, 1–16.

9. Newell, K. M., 1986, Constraints on the development of coordination, in M. G. Wade and H. T. A. Whiting, eds., *Motor Development in Children: Aspects of Coordination and Control*, 341–361, Amsterdam: Martinus Nijhoff.

10. Lewin, K., 1936, *Principles of Topological Psychology*, New York: McGraw-Hill.

11. Kirlik, A., 2004, On Stoffregen's definition of affordance, *Ecological Psychology.*, *16*, 73–77.

12. Abbate, A. J., and Bass, E. J., 2017, Modeling affordance using formal methods, *Proceedings of the Human Factors and Ergonomics Society 2017 Annual Meeting*, 723–727, Santa Monica, CA: Human Factors and Ergonomics Society.

13. Abbate and Bass, 2017.

14. Fitts, P. M., and Deininger, R. L., 1954, S-R compatibility: Correspondence among paired elements within stimulus and response codes, *Journal of Experimental Psychology*, *48*, 483–492; Fitts, P. M., and Seeger, C. M., 1953, S-R compatibility: Spatial characteristics of stimulus and response codes, *Journal of Experimental Psychology*, *46*, 199–210; Welford, A. T., 1976, *Skilled Performance. Perceptual and Motor Skills*, Glenview, IL: Foreman.

15. Kornblum, S., Hasbroucq, T., and Osman, A., 1990, Dimensional overlap: Cognitive basis for stimulus-response compatibility—A model and taxonomy, *Psychological Review*, *97*, 253; Kornblum, S., and Lee, J. W., 1995, Stimulus-response compatibility with relevant and irrelevant stimulus dimensions that do and do not overlap with the response, *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 855.

16. Proctor, R. W., and Vu, K. P. L., 2016, Principles for designing interfaces compatible with human information processing, *International Journal of Human-Computer Interaction*, *32*, 2–22.

17. Fitts and Seeger, 1053.

18. Stins, J. F., and Michaels, C. F., 1997, Stimulus-response compatibility is information-action compatibility.,*Ecological Psychology, 9*, 25–45.

19. Norman, D. A., 1999, Affordance, conventions, and design, *Interactions*, *6*, 38–43.

20. Turner, P., 2005, Affordance as context, *Interacting with Computers*, *17*, 787–800.

21. Stefanucci, J. K., and Geuss, M. N., 2009, Big people, little world: The body influences size perception, *Perception, 38*, 1782–1795; Warren, W. H., and Wang, S., 1987, Visual guidance of walking through apertures: Body-scaled information for affordances, *Journal of Experimental Psychology: Human Perception and Performance*, *13*, 371–383.

22. Warren, W. H., 1984, Perceiving affordances: Visual guidance of stair climbing, *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 683–703.

23. Kinsella-Shaw, J. M., Shaw, B., and Turvey, M., 1992, Perceiving walk-on-able slopes, *Ecological Psychology*, *4*, 223–239.

24. Mark, L. S., 1987, Eye-height-scaled information about affordances: A study of sitting and stair climbing, *Journal of Experimental Psychology: Human Perception and Performance*, *13*, 361–370.

25. Cornus, S., Montagne, G., and Laurent, M., 1999, Perception of a stepping-across affordance, *Ecological Psychology*, *11*, 249–267.

26. Stefanucci, J. K., and Geuss, M. N., 2010, Duck! Scaling the height of a horizontal barrier to body height, *Attention, Perception and Psychophysics*, *72*, 1338–1349.

27. Franchak, J. M., and Adolph, K. E., 2014, Gut estimates: Pregnant women adapt to changing possibilities for squeezing through doorways, *Attention, Perception, & Psychophysics*, *76*, 460–472.

28. Higuchi, T., Murai, G., Kijima, A., Seya, Y., Wagman, J. B., and Imanaka, K., 2011, Athletic experience influences shoulder rotations when running through apertures, *Human Movement Science*, *30*, 534–549.

29. Profitt, D. R., 2006, Embodied perception and the economy of action, *Perspectives on Psychological Science*, *1*, 110–122.

30. Proffit, D. R., and Linkenauger, S. A., 2013, Perception viewed as a phenotypic expression, in W. Prinz, M. Biesert, and A. Herwig, eds., *Action Science: Foundations of an Emerging Discipline*, 171–179, Cambridge, MA: MIT Press.

31. Witt, K. L., and Proffit, D. R., 2008, Action-specific influences on distance perception: A role for motor simulation, *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 1479–1492.

32. Humphreys, G. W., 2001, Objects, affordances, action, *Psychologist*, *14*, 408–412.

33. Malafouris, L., 2013, *How Things Shape the Mind: A Theory of Material Engagement*, Cambridge, MA: MIT Press.

34. Kiverstein, J., and Rietveld, E., 2020, Scaling-up skilled intentionality to linguistic thought, *Synthese*, 198, 175–194.

35. Bruineberg, J., and Rietveld, E., 2014, Self-organisation, free-energy minimisation, and optimal grip on field of affordances, *Frontiers in Human Neuroscience, https://doi.org/10.3389/fnhum.2014.00599;* Rietveld, E., and Kiverstein, J., 2014, A rich landscape of affordances, *Ecological Psychology*, *26*, 325–352.

36. Rietveld, E., De Haan, S., and Denys, D., 2013, Social affordances in context: What is it that we are bodily responsive to? Commentary on Schilbach et al., *Behavioral and Brain Sciences*, *36*, 436.

37. Rietveld and Kiverstein, 2014, 335.

38. Rietveld, E., Denys, D., and Van Westen, M., 2018, Ecological-enactive cognition as engaging with a rich landscape of affordances: The skilled intentionality framework (SIF), in A. Newen, L. De Bruin, and S. Gallagher, eds., *The Oxford handbook of 4E cognition*, New York: Oxford University Press, 41–70.

39. Ingold, T., 2000, *The Perception of the Environment: Essays on Livelihood, Dwelling and Skill*, London: Psychology Press, 5.

40. Käufer, S., and Chemero, A., 2015, *Phenomenology: An Introduction*, Cambridge: Polity Press.

41. Rietveld and Kiverstein, 2014.

42. Wittgenstein, L., 1953, *Philosophical Investigations*, Oxford: Blackwell.

43. Gibson, J. J., 1966, *The Senses Considered as Perceptual Systems*, Boston: Houghton Mifflin.

44. Gibson, J. J., 1979, *The Ecological Approach to Visual Perception*, Boston: Houghton Mifflin, 127.

45. Chemero, A., 2003, An outline of a theory of affordances, *Ecological Psychology*, *15*, 181–195.

46. Gibson, 1979, 128.

47. Pickering, A., 1995, *The Mangle of Practice: Time, Agency, and Science*, Chicago: University of Chicago Press.

48. Davis, J., 2020, *How Artifacts Afford: The Power and Politics of Everyday Things*, Cambridge, MA: MIT Press.

49. Merleau-Ponty, M., 2014 [1945], *Phenomenology of Perception*, London: Routledge

50. Davis, 2020, 97.

51. Koffka, K., 2001 [1935], *Principles of Gestalt Psychology*, London: Routledge, 353.

52. Koffka, 2001, 138.

53. Käufer and Chemero, 2015.

54. Koffka, 1935, 7. The quotation ends with the phrase "*and a woman says 'Love me.'*" The contemporary reader will be affronted by the misogyny here. One might have an idea of what the "woman" in question might "say" to Koffka (although I doubt that it would be printable).

55. Peirce, C. S., 1991, *Peirce on Signs: Writings on Semiotic*, J. Hoopes, ed., Chapel Hill: University of North Carolina.

56. Gaver, W., 1991, *Technology Affordances, CHI"91*, New York: Association for Computing Machinery.

57. Lewin,1936.

58. Gibson, 1979, 138–139.

59. Gibson, 1979, 138–139.

60. Hone, K. S., and Baber, C., 2001, Designing habitable dialogues for speech-based interaction with computers, *International Journal of Human-Computer Studies*, *54*, 637–662.

61. Chalmers, M. and Galani, A., 2004, Seamful interweaving: Heterogeneity in the theory and design of interactive systems, *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, 243–252, Cambridge, MA: Association for Computing Machinery.

Chapter 5

1. Carroll, J. M., and Rosson, M. B., 1992, Getting around the task-artifact cycle: How to make claims and design by scenario, *ACM Transactions on Information Systems (TOIS)*, *10*, 181–212.

2. Bijker, W. E., 1997, *Of Bicycles, Bakelites, and Bulbs: Toward a Theory of Sociotechnical Change*, Cambridge, MA: MIT Press.

3. Tuomi, I., 2002, *Networks of Innovation: Change and Meaning in the Age of the Internet*, Oxford: Oxford University Press.

4. Rasmussen, J., 1997, Risk management in a dynamic society: A modelling problem, *Safety science*, *27*, 183–213.

5. Eason, K., 2001, Changing perspectives on the organizational consequences of information technology, *Behaviour & Information Technology*, *20*, 323–328.; Klein, L., 2014, What do we actually mean by “sociotechnical”? On values, boundaries and the problems of language, *Applied Ergonomics*, *45*, 137–142.

6. Latour, B., 2013, Reassembling the social: An introduction to actor-network-theory, *Journal of Economic Sociology*, *14*, 73–87.

7. Baber, C., Morar, N. S., and McCabe, F., Ecological interface design, the proximity compatibility principle, and automation reliability in road traffic management, *IEEE Transactions on Human-Machine Systems*, *49*, 241–249.

8. Burns, C. M., and Hajdukiewicz, J., 2004, *Ecological Interface Design*, San Diego, CA: CRC Press; Jenkins, D. P., Stanton, N. A., Salmon, P. M., and Walker, G. H., 2008, *Cognitive Work Analysis: Coping with Complexity*, Avebury: Ashgate; Vicente, K. J., and Rasmussen, J., 1992, Ecological interface design: Theoretical foundations, *IEEE Transactions on Systems, Man and Cybernetics*, *22*, 589–606.

9. Baber, C., Attfield, S., Conway, G., Rooney, C., and Kodagoda, N., 2016, Collaborative sensemaking during intelligence analysis exercises, *International Journal of Human-Computer Studies*, *86*, 94–108.

10. Lintern, G., 2010, A comparison of the decision ladder and the recognition-primed decision model, *Journal of Cognitive Engineering and Decision Making*, *4*, 304–327.

11. Rasmussen, J., 1986, A framework for cognitive task analysis in systems design, in E. Hollnagel, E. Mancini, and D. Woods, eds., *Intelligent Decision Support in Process Environments*, 175–196, Berlin: Springer.

12. Rasmussen, 1986.

13. Mayes, J. T., Draper, S. W., McGregor, A. M., and Oatley, K., 1988, Information flow in a user interface: The effect of experience and context on the recall of MacWrite screens, in D. M. Jones and R. Winder, eds., *Human-Computer Interaction*, 275–289, Cambridge: Cambridge University Press.

14. Naikar, N., 2010, *A Comparison of the Decision Ladder Template and the Recognition-Primed Decision Model*, Defence Science and Technology Organisation, Victoria (Australia), Air Operations Division.

15. Lintern, 2010.

16. Bennett, K. B., and Flach, J., 2019, Ecological interface design: Thirty-plus years of refinement, progress, and potential, *Human Factors*, *61*, 513–525; Bennett, K. B.,

and Flach, J. M., 2011, *Display and Interface Design: Subtle Science, Exact Art*, Boca Raton, FL: CRC Press; Burns, C. M. and Hajdukiewicz, J., 2004, Ecological Interface Design, Boca Raton, FL: CRC Press.

17. Burns, C. M., 2000, Navigation strategies with ecological displays, *International Journal of Human-Computer Studies*, *52*, 111–129.

18. Vicente, K. J., and Rasmussen, J., 1992, Ecological interface design: Theoretical foundations, *IEEE Transactions on Systems, Man and Cybernetics*, *22*, 589–606.

19. Rasmussen, J., 1999, Ecological interface design for reliable human-machine systems, *International Journal of Aviation Psychology*, *9*, 203–223.

20. Vicente and Rasmussen, 1992.

21. Lau, N., and Jamieson, G. A., 2006, Ecological interface design for the condenser subsystems of a boiling water reactor simulator, in E. Koningsveld, *Proceedings of the 16th World Congress on Ergonomics: IEA 2006*, Maastricht, The Netherlands: Elsevier.

22. Vicente, K. J., and Rasmussen, J., 1988, On applying the skills, rules, knowledge framework to interface design, *Proceedings of the Human Factors Society Annual Meeting*, Los Angeles, CA: SAGE Publications, 254–258.

23. Adapted from Vicente, K. J., 1999, *Cognitive Work Analysis*, Mahwah, NJ: Erlbaum.

24. Adapted from Vincente, 1999.

25. Adapted from Vicente, 1999.

26. Bennett, K. B., Posey, S. M., and Shattuck, L. G., 2008, Ecological interface design for military command and control, *Journal of Cognitive Engineering and Decision Making*, *2*, 349–385.

27. Gibson, J. J., 1979, *The Ecological Approach to Visual Perception*, Boston, Houghton Mifflin, 37.

28. Duez, P., and Vicente, K. J., 2005, Ecological interface design and computer network management: The effects of network size and fault frequency, *International journal of human-computer studies*, *63*, 565–586; Ham, D. H., and Yoon, W. C., 2001a, Design of information content and layout for process control based on goal–means domain analysis, *Cognition, Technology & Work*, *3*, 205–223; Ham, D. H., and Yoon, W. C., 2001b, The effects of presenting functionally abstracted information in fault diagnosis tasks, *Reliability Engineering & System Safety*, *73*, 103–119; Lau, N., Jamieson, G. A., Skraaning Jr., G., and Burns, C. M., 2008, Ecological interface design in the nuclear domain: An empirical evaluation of ecological displays for the secondary subsystems of a boiling water reactor plant simulator, *IEEE Transactions on Nuclear Science*, *55*, 3597–3610; Pawlak, W. S., and Vicente, K. J., 1996, Inducing effective operator control through ecological interface design, *International Journal of Human-Computer Studies*, *44*, 653–688.

29. Carrasco, C. G. A., and Jamieson, O. St-cyr., 2014, Revisiting three ecological interface design experiments to investigate performance and control stability effects under normal conditions, IEEE International Conference on Systems, Man and Cybernetics, New York: IEEE, 323–328.

30. Burns, 2000; Ham and Yoon, 2001a; Ham et al., 2008, ibid; Janzen, M. E., and Vicente, K. J., 1998, Attention allocation within the abstraction hierarchy, *International Journal of Human-Computer Studies*, *48*, 521–545.

31. Borst, C., Flach, J. M., and Ellerbroek, J., 2014, Beyond ecological interface design: Lessons from concerns and misconceptions, *IEEE Transactions on Human-Machine Systems*, *45*, 164–175.

32. Vicente and Rasmussen, 1992.

Chapter 6

1. Dourish, P., 2004, *Where the Action Is: The Foundations of Embodied Interaction*, Cambridge, MA: MIT Press, 126.

2. Ishii, H., Mazalek, A., and Lee, J., 2001, Bottles as a minimal interface to access digital information, in *CHI"01 Extended Abstracts on Human Factors in Computing Systems*, 187–188, New York: Association of Computing Machinery.

3. Underkoffler, J., and Ishii, H., 1999, Urp: A luminous-tangible workbench for urban planning and design, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 386–393, New York: Association of Computing Machinery.

4. Rekimoto, J., 1997, Pick-and-drop: A direct manipulation technique for multiple computer environments, in G. Robertson and C. Schmandt, eds., *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*, 31–39, New York: Association of Computing Machinery.

5. Levy, S., 1994, *Insanely Great: The Life and Times of Macintosh, the Computer That Changed Everything*, London: Penguin.

6. CHI 2019 SIGCHI Lifetime Research Award, Hiroshi Ishii: Making Digital Tangible, https://www.youtube.com/watch?v=T6D1A9J0qIE.

7. Winograd, T., Flores, F., and Flores, F. F., 1986, *Understanding computers and cognition: A New Foundation for Design*, Norwood, NJ: Ablex Publishing Corporations.

8. Weiser, M., 2002, The computer for the 21st century, *IEEE Pervasive Computing*, *1*, 19–25.

9. Dreyfus, H., 1991, *What Computers Still Can"t Do*, Cambridge, MA: MIT Press.

10. Dotov, D. G., Nie, L., and Chemero, A., 2010, A demonstration of the transition from ready-to-hand to unready-to-hand, *PLoS One*, *5*, e9433.

11. Dourish, 2001.

12. Niedenthal, P. M., Barsalou, L. W., Winkielman, P., Krauth-Gruber, S., and Ric, F., 2005, Embodiment in attitudes, social perception, and emotion, *Personality and Social Psychology Review*, *9*, 184–211.

13. Bianchi-Berthouze, N., Kim, W. W., and Patel, D., 2007, Does body movement engage you more in digital game play? And why? *International Conference on Affective Computing and Intelligent Interaction*, 102–113, Berlin: Springer.

14. Antle, A. N., Corness, G., and Droumeva, M., 2009, What the body knows: Exploring the benefits of embodied metaphors in hybrid physical digital environments, *Interacting with Computers*, *21*, 66–75.

15. Hurtienne, J., and Blessing, L., 2007, Design for Intuitive Use—Testing image schema theory for user interface design, in *DS 42: Proceedings of ICED 2007, the 16th International Conference on Engineering Design, Paris, France, 28.-31.07. 2007*, Glasgow: The Design Society, 829–830; Hurtienne, J., Israel, J. H., and Weber, K., 2008, Cooking up real world business applications combining physicality, digitality, and image schemas, *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction*, New York: Association of Computing Machinery, 239–246.

16. Fällman, D., 2003, In romance with the materials of mobile interaction: A phenomenological approach to the design of mobile information technology , PhD dissertation, Umeå University, Faculty of Social Sciences, Informatic.

17. Hornecker, E., 2005, A design theme for tangible interaction: Embodied facilitation, in *ECSCW 2005: Proceedings of the 16th European Conference on Computer-Supported Cooperative Work: The International Venue on Practice-centred Computing and the Design of Cooperation Technologies - Exploratory Papers, Reports of the European Society for Socially Embedded Technologies*, 23–43, Springer, Dordrecht.

18. Marshall, P., and Hornecker, E., 2013, Theories of embodiment in HCI, in *The SAGE Handbook of Digital Technology Research*, 1:144–158, New York: Sage.

19. Stoffregen, T. A., Bardy, B. G., and Mantel, B., 2006, Affordances in the design of enactive systems, *Virtual Reality*, *10*, 4–10.

20. Baber, C., Khattab, A., Russell, M., Hermsdörfer, J., and Wing, A., 2017, Creating affording situations: Coaching through animate objects, *Sensors*, *17*, 2308.

21. Daniel Leithinger, Sean Follmer, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii, n.d., inFORM, MIT Media Lab, Tangible Media Group, https://tangible.media.mit.edu/project/inform/.

22. Brooks, R., 1991, Intelligence without representation, *Artificial Intelligence Journal*, *47*, 139–160.

23. Agre, P. E., 1991, *The Dynamic Structure of Everyday Life*, Cambridge: Cambridge University Press; Agre, P. E., and Chapman, D., 1990, What are plans for? in P. Maes, ed.,

Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back, 17–34, Cambridge, MA: MIT Press.

24. Brooks, 1991, 15.

25. Brooks, 1991, 16.

26. Brooks, 1991, 16.

27. Brooks, 1991, 16.

28. "Furby HBW" by Vox Efx is licensed under CC BY 2.0.

29. Negroponte, N., 1995, *Being Digital*, New York: Alfred A Knopf.

30. Dennett, D.C., 1989, *The Intentional Stance,* Cambridge, MA: MIT Press.

31. Weiser, M., 1999, The computer for the 21st century, *Scientific American*, *265*, 94–105.

32. Society of Automotive Engineers, 2016, Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles (Surface Vehicle Recommended Practice: Superseding J3016 Jan 2014), SAE International, September.

33. Sheridan T. B., and Verplank, W. L., 1978, Human and computer control of undersea teleoperators (Tech. Rep.). Cambridge, MA: Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT.

34. Bainbridge, L., 1983, Ironies of automation, *Automatica*, *19*, 775.

35. Licklider, J. C., 1960, Man-computer symbiosis, *IRE Transactions on Human Factors in Electronics*, *1*, 4–11.

36. Latour, B., 2013, Reassembling the social: An introduction to actor-network-theory, *Journal of Economic Sociology*, *14*, 73–87.

37. Latour, 2013.

38. Kätsyri, J., Förger, K., Mäkäräinen, M., and Takala, T., 2015, A review of empirical evidence on different uncanny valley hypotheses: Support for perceptual mismatch as one road to the valley of eeriness, *Frontiers in Psychology*, *6*, 390.

39. Heidegger, M., 1962, *Being and Time*, New York: Harper Row.

40. Latour, B., 1999, *Pandora'sHope, An Essay on the Reality of Science Studies*, Cambridge, MA: Harvard University Press, 181.

41. Pickering, A., 1995, *The Mangle of Practice: Time, Agency, and Science*, Chicago: University of Chicago Press.

42. "File: Lexan bubble chamber.jpg" by CERN is licensed under CC BY-SA 4.0.

43. Maturana, H., 1987, Everything is said by an observer, in W. I. Thompson, ed., *GAIA: A Way of Knowing*, Hudson, NY: Lindisfarne Press, 73.

44. Bång, M., and Timpka, T., 2003, Cognitive tools in medical teamwork: The spatial arrangement of patient records, *Methods of information in medicine*, *42*, 331–336; Nemeth, C., 2003, How cognitive artifacts support distributed cognition in acute care, *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society*, 381–385, Santa Monica, CA: Human Factors and Ergonomics Society; Seagull, F. J., Plasters, C., Xiao, Y., and Mackenzie, C. F., 2003, Collaborative management of complex coordination systems: Operating room schedule coordination, *Proceedings of the 47th Annual Meeting of The Human Factors and Ergonomics Society*, 1521–1525, Santa Monica, CA: Human Factors and Ergonomics Society.

45. Scaife, M., and Rogers, Y., 1996, External cognition: How do graphical representations work? *International Journal of human-Computer Studies*, *45*, 185–213.

Chapter 7

1. Still, J. D., and Dark, V. J., 2013, Cognitively describing and designing affordances, *Design Studies*, *34*, 285–301.

2. Tucker, M., and Ellis, R., 1998, On the relations of seen objects and components of potential actions, *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 830–846.

3. Craighero, L., Fadiga, L., Rizzolatti, G., and Umiltà, C., 1999, Action for perception: A motor-visual attentional effect, *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1673; Ellis, R., and Tucker, M., 2000, Micro-affordance: The potentiation of components of action by seen objects, *British Journal of Psychology*, *91*, 451–471; Klatzky, R. L., Fikes, T. G., and Pellegrino, J. W., 1995, Planning for hand shape and arm transport when reaching for objects, *Acta Psychologica*, *88*, 209–232; Riddoch, M. J., Edwards, M. G., Humphreys, G. W., West, R., and Heafield, T., 1998, Visual affordances direct action: Neuropsychological evidence from manual interference, *Cognitive Neuropsychology*, *15*, 645–683; Tucker and Ellis, 1998; Tucker, M., and Ellis, R., 2001, The potentiation of grasp types during visual object categorization, *Visual Cognition*, *8*, 769–800.

4. Tipper, S. P., Howard, L. A., and Jackson, S. R., 1997, Selective reaching to grasp: Evidence for distractor interference effects, *Visual Cognition*, *4*, 1–38.

5. Tipper, S. P., Howard, L. A., and Houghton, G., 1998, Action–based mechanisms of attention, *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *353*, 1385–1393.

6. Rosenbaum, D. D., Chapman, K. M., Weigelt, M., Weiss, D. J., and van der Wel, R., 2012, Cognition, action, and object manipulation, *Psychological Bulletin*, *138*, 924–946.

7. Chapman, S., 1968, Catching a baseball, *American Journal of Physics*, *36*, 868–870; Fink, P., Foo, P., and Warren, W., 2009, Catching fly balls in virtual reality: A critical test of the outfielder problem, *Journal of Vision*, *9*, 14.

8. Saxberg, B. V., 1987, Projected free fall trajectories, *Biological Cybernetics*, *56*, 159–175; Saxberg, B. V., 1987, Projected free fall trajectories. II. Human experiments. *Biological Cybernetics*, *56*, 177–184.

9. Linear optical trajectory (LOC): Mcbeath, M. K., Shaffer, D. M., and Kaiser, M. K., 1995, How baseball outfielders determine where to run to catch fly balls, *Science*, *268*, 569–573.

10. Optical acceleration cancellation (OAC): Chapman, 1968.

11. Chemero, A., 2009, *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press, 191.

12. Haken, H., Kelso, J. S., and Bunz, H., 1985, A theoretical model of phase transitions in human hand movements, *Biological Cybernetics*, *51*, 347–356.

13. Craik, K. J., 1947, Theory of the human operator in control systems. I. The operator as an engineering system, *British Journal of Psychology. General Section*, *38*, 56–61; Craik, K. J., 1948, Theory of the human operator in control systems. II. Man as an element in a control system, *British Journal of Psychology*, *38*, 142.

14. Kugler, P.N., Kelso, J. A.S., and Turvey, M. T., 1982, On the control and coordination of naturally developing systems, in J. A. S. Kelso and J. E. Clark, eds., *The Development of Movement Control and Coordination*, 25–36, New York: Wiley; Kugler, P.N., and Turvey, M. T., 1987, *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*, Hillsdale, NJ: Lawrence Erlbaum Associates.

15. Haken H., 1983, *Synergetics, an Introduction: Non-Equilibrium Phase Transitions and Self-Organisation in Physics, Chemistry and Biology*, Berlin: Springer-Verlag.

16. Biryukova, E., and Bril, B., 2012, Biomechanical analysis of tool use: A return to Bernstein's tradition, *Zeitschrift für Psychologie*, *220*, 53.

17. Baber, C., Cengiz, T. G., Starke, S., and Parekh, M., 2015, Objective classification of performance in the use of a piercing saw in jewellery making, *Applied Ergonomics*, *51*, 211–221.

18. Findlay, J. M., and Gilchrist, I. D., 2003, *Active Vision: The Psychology of Looking and Seeing*, Oxford: Oxford University Press.

19. Chemero, A., 2020, Epilogue: What embodiment is, in N. K. Dees, ed., *A Multidisciplinary Approach to Embodiment*, 133–139, London: Routledge.

20. Baber, C., Chen, X., and Howes, A., 2015, (Very) rapid decision-making: Framing or filtering? , International Conference on Naturalistic Decision-Making, Mc-Lean, VA.

21. James, W., 1890, *The Principles of Psychology*, vol. 1, New York: Holt.

22. Knill, D. C., and Pouget, A., 2004, The Bayesian brain: The role of uncertainty in neural coding and computation, *Trends in Neurosciences*, *27*, 712–719.

23. Moran, R. J., Campo, P., Symmonds, M., Stephan, K. E., Dolan, R. J., and Friston, K. J., 2013, Free energy, precision and learning: The role of cholinergic neuromodulation, *Journal of Neuroscience*, 33, 8227–8236.

24. Clark, A., 2015, *Embodied Prediction*, Frankfurt: MIND Group.

25. Friston, K., 2010, The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138; Friston, K., Kilner, J., and Harrison, L., 2006, A free energy principle for the brain, *Journal of Physiology—Paris*, *100*, 70–87.

26. Friston, 2010.

27. Seth, A. K., 2015, The cybernetic Bayesian brain—From interoceptive inference to sensorimotor contingencies, in T. Metzinger and J. M. Windt, eds, *Open MIND*: 1–24, Frankfurt am Main: MIND Group.

28. O'Regan, K., and Noe, A., 2001 A sensorimotor account of vision and visual consciousness, *Behavioral Brain Science*, *24*, 939–1031.

29. Clark, A., Tower, D. H., and Jx, E., 2006, Vision as dance? The challenges for sensorimotor contingency theory, *Psyche*, *12*, 1–10.

30. Kirsh, D., 2013, Embodied cognition and the magical future of interaction design, *ACM Transactions on Computer-Human Interaction (TOCHI)*, *20*, 10.

31. Ingold, T., 2010, The textility of making, *Cambridge Journal of Economics*, *34*, 91–10.

32. Kloos, H., and Van Orden, G. C., 2009, Soft-assembled mechanisms for the unified theory, in J. Spencer, ed., *Toward a Unified Theory of Development: Connectionism and Dynamic Systems Theory Re-Considered*, 259, Oxford: Oxford University Press.

33. Richardson, M. J., and Chemero, A., 2014, Complex dynamical systems and embodiment, in L. Shapiro, ed., *The Routledge Handbook of Embodied Cognition*, 39–50, London: Routledge.

34. Dotov, D. G., Nie, L., and Chemero, A., 2010, A demonstration of the transition from ready-to-hand to unready-to-hand, *PLoS One*, *5*, e9433.

35. Baber, C., and Starke, S. D., 2015, Using 1/f scaling to study variability and dexterity in simple tool-using tasks, *Proceedings of the 59th Human Factors and Ergonomics Society Annual Meeting*, 431–435, Santa Monica, CA: HFES.

36. Van Orden, G. C., Holden, J. G., and Turvey, M. T., 2005, Human cognition and 1/f scaling, *Journal of Experimental Psychology: General*, *134*, 117.

37. Kello, C. T., Beltz, B. C., Holden, J. G., and Van Orden, G. C., 2007, The emergent coordination of cognitive function, *Journal of Experimental Psychology: General*, *136*, 551.

38. Wiltshire, T. J., Steffensen, S. V., and Fiore, S. M., 2019, Multiscale movement coordination dynamics in collaborative team problem solving, *Applied Ergonomics*, *79*.

Chapter 8

1. Shaban, D., and Koehler, L., 2016, *The Eames Creative Process*, YouTube, watch?v=tUOFhnWTbm0.

2. "Eames Chair" by designsbykari is licensed under CC BY 2.0.

3. Glaveanu, V. P., 2014, *Distributed Creativity: Thinking outside the Box of the Creative Individual*. Cham, Germany: Springer International Publishing.

4. Woodman, R. W., and Schoenfeldt, T., 1989, Individual differences in creativity: An interactionist perspective, in J. A. Glover, R. R. Ronning, and C. R. Reynolds, eds., *Handbook of Creativity*, 77–93, New York: Plenum Press.

5. Leschziner, V., and Brett, G., 2019, Beyond two minds: cognitive, embodied, and evaluative processes in creativity, *Social Psychology Quarterly, 82*, 340–366.

6. Sawyer, R. K., and de Zutter, S., 2009, Distributed creativity: How collective creations emerge from collaboration, *Psychology of Aesthetics, Creativity, and the Arts*, *3*, *81–92*.

7. Marx, K., 1990 [1886], *Capital, Volume 1*, London: Penguin Classics, 283.

8. Ingold, T., 2013, *Making: Anthropology, Archaeology, Art and Architecture*, New York: Routledge.

9. Pallasmaa, J., 2009. *The Thinking Hand: Existential and Embodied Wisdom in Architecture*, Chichester: Wiley.

10. Ingold, 2013.

11. Runeson, S., 1989, A note on the utility of ecologically incomplete invariants, *Newsletter of the International Society for Ecological Psychology*, *4*, 6–9.

12. Baber, C., 2006, Cognitive aspects of tool use, *Applied Ergonomics*, *37*, 3–15.

13. Rosenbaum, D. D., Chapman, K. M., Weigelt, M., Weiss, D. J., and van der Wel, R., 2012, Cognition, action, and object manipulation, *Psychological Bulletin*, *138*, 924–946.

14. Ramstead, M. J. D., Veissiere, S. P. L., and Kirmayer, L. J., 2016, Cultural affordances: Scaffolding local worlds through shared intentionality and regimes of attention, *Frontiers in Psychology*, *7*, 1090.

15. Davis, J., 2020, *How Artifacts Afford: The Power and Politics of Everyday Things*, Cambridge, MA: MIT Press.

16. Bardram, J. E., 2005, Activity-based computing: Support for mobility and collaboration in ubiquitous computing, *Personal and Ubiquitous Computing*, *9*, 312–322.

17. Dourish, P., 2004, *Where The Action Is: The Foundations of Embodied Interaction*, Cambridge, MA: MIT Press, 126.

18. Dourish, 2004, 126.

19. Dourish, 2004, 138.

20. Csikszentmihályi, M., 1990, The domain of creativity, in M. A. Runco and R. S. Albert, eds., *Sage Focus Editions, Vol. 115. Theories of Creativity*, 190–212, London: Sage.

21. Withagen, R., and van der Kamp, J., 2018, An ecological approach to creativity in making, *New Ideas in Psychology*, *49*, 1–6.

Index