

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-Tourangeau² 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 of 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 count⁶ (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 it 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

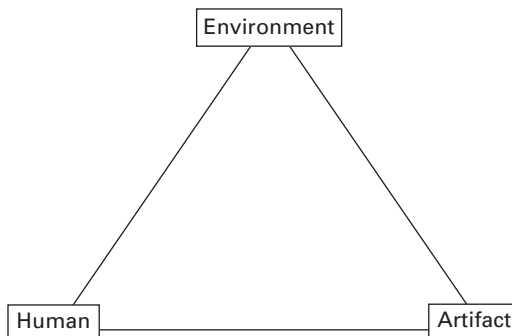


Figure 1.1
Interacting elements of a human-artifact-environment system.

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 memory¹³ 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 1956¹⁴). 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 operators²⁰). But, this is markedly different from the idea of a mental model that, say, Frederic Bartlett²¹ 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), Shannon³⁹ 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 cognitive⁴⁴. 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 cognition⁴⁵ 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, Hutchins⁴⁷ 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, Nemeth⁴⁸ 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 Chalmers⁴⁹ 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 Aizawa⁵⁰ 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 neurons⁵⁷). 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, Bernstein⁶⁰ 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 musculo-skeletal 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 Johnson⁶⁸ 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 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."

