

Six Views of Embodied Cognition

Margaret Wilson
University of California, Santa Cruz

Address correspondence to:
Margaret Wilson
Department of Psychology
University of California, Santa Cruz
Santa Cruz CA
mlwilson@cats.ucsc.edu

IN PRESS: PSYCHONOMIC BULLETIN & REVIEW

Abstract

The emerging viewpoint of embodied cognition holds that cognitive processes are deeply rooted in the body's interactions with the world. This position actually houses a number of distinct claims, some of which are more controversial than others. This paper distinguishes and evaluates the following six claims: 1) cognition is situated; 2) cognition is time-pressured; 3) we off-load cognitive work onto the environment; 4) the environment is part of the cognitive system; 5) cognition is for action; 6) off-line cognition is body-based. Of these, the first three and the fifth appear to be at least partially true, and their usefulness is best evaluated in terms of the range of their applicability. The fourth claim, I argue, is deeply problematic. The sixth claim has received the least attention in the literature on embodied cognition, but it may in fact be the best documented and most powerful of the six claims.

There is a movement afoot in cognitive science to grant the body a central role in shaping the mind. Proponents of embodied cognition take as their theoretical starting point, not a mind working on abstract problems, but a body that requires a mind to make it function. These opening lines by Clark (1998) are typical: “Biological brains are first and foremost the control systems for biological bodies. Biological bodies move and act in rich real-world surroundings.”

Traditionally, the various branches of cognitive science have viewed the mind as an abstract information processor, whose connections to the outside world were of little theoretical importance. Perceptual and motor systems, though reasonable objects of inquiry in their own right, were not considered relevant to understanding “central” cognitive processes. Instead, they were thought to serve merely as peripheral input and output devices. This stance was evident in the early decades of cognitive psychology, when most theories of human thinking dealt in propositional forms of knowledge. During the same time period, artificial intelligence was dominated by computer models of abstract symbol processing. Philosophy of mind, too, made its contribution to this zeitgeist, most notably in Fodor’s modularity hypothesis (1983). According to Fodor, central cognition is not modular, but its connections to the world are. Perceptual and motor processing are done by informationally encapsulated plug-ins providing sharply limited forms of input and output.

However, there is a radically different stance that also has roots in diverse branches of cognitive science. This stance has emphasized sensory and motor functions, and their importance for successful interaction with the environment. Early sources include the view of 19th century psychologists that there was no such thing as “imageless thought” (Goodwin, 1999); motor theories of perception such as those suggested by William James and others (see Prinz, 1987, for review); the developmental psychology of Jean Piaget, which emphasized the emergence of cognitive abilities out of a groundwork of sensorimotor abilities; and the ecological psychology of J. J. Gibson, which viewed perception in terms of *affordances* – potential interactions with the environment. In the 1980’s, linguists began exploring how abstract concepts may be based on metaphors for bodily, physical concepts (e.g. Lakoff & Johnson, 1980). At the same time, within the field of artificial intelligence, behavior-based robotics began to emphasize routines for interacting with the environment rather than internal representations used for abstract thought (e.g. Brooks, 1986).

This kind of approach has recently attained high visibility, under the banner of embodied cognition. There is a growing commitment to the idea that the mind must be understood in the context of its relationship to a physical body that interacts with the world. It is argued that we evolved from creatures whose neural resources were devoted primarily to perceptual and motoric processing, and whose cognitive activity consisted largely of immediate, on-line interaction with the environment. Hence human cognition, rather than being centralized, abstract, and sharply distinct from peripheral input and output modules, may instead have deep roots in sensorimotor processing.

While this general approach is enjoying increasingly broad support, there is in fact a great deal of diversity in the claims involved and the degree of controversy they attract. If the term “embodied cognition” is to retain meaningful use, we need to disentangle and evaluate these diverse claims. Among the most prominent are the following:

- 1) Cognition is situated. Cognitive activity take place in the context of a real-world environment, and inherently involves perception and action.
- 2) Cognition is time-pressured. We are “mind on the hoof” (Clark, 1997), and cognition must be understood in terms of how it functions under the pressures of real-time interaction with the environment.
- 3) We off-load cognitive work onto the environment. Because of limits on our information-processing abilities (e.g. limits on attention and working memory), we exploit the environment to reduce the cognitive workload. We make the environment hold or even manipulate information for us, and we harvest that information only on a need-to-know basis.
- 4) The environment is part of the cognitive system. The information flow between mind and world is so dense and continuous that, for scientists studying the nature of cognitive activity, the mind alone is not a meaningful unit of analysis.
- 5) Cognition is for action. The function of the mind is to guide action, and cognitive mechanisms such

as perception and memory must be understood in terms of their ultimate contribution to situation-appropriate behavior.

6) Off-line cognition is body-based. Even when decoupled from the environment, the activity of the mind is grounded in mechanisms that evolved for interaction with the environment – that is, mechanisms of sensory processing and motor control.

Frequently in the literature on embodied cognition several or all of these claims are presented together as if they represented a single point of view. This strategy may have its uses, for example in helping to draw a compelling picture of what embodied cognition might be and why it might be important. This may have been particularly appropriate at the time that attention first was drawn to this set of ideas, when audiences were as yet unfamiliar with this way of conceptualizing cognition. The time has come, though, to take a more careful look at each of these claims on its own merits.

Claim 1: Cognition is Situated

A cornerstone of the embodied cognition literature is the claim that cognition is a situated activity (e.g. Chiel & Beer, 1997; Clark, 1997; Pfeifer & Scheier, 1999; Steels & Brooks, 1995; a commitment to situated cognition can also be found in the literature on dynamical systems, e.g. Beer, 2000; Port & van Gelder, 1995; Thelen & Smith, 1994; Wiles & Dartnall, 1999). Some authors go so far as to complain that the phrase “situated cognition” implies, falsely, that there also exists cognition that is not situated (Greeno & Moore, 1993, p. 50). It is important, then, that we be clear on what exactly it means for cognition to be situated.

Simply put, situated cognition is cognition that takes place in the context of task-relevant inputs and outputs. That is, while a cognitive process is being carried out, perceptual information continues to come in that affects processing, and motor activity is executed that affects the environment in task-relevant ways. Driving, hold a conversation, and moving around a room while trying to imagine where the furniture should go, are all cognitive activities that are situated in this sense.

Even with this basic definition what it means for cognition to be situated, we can note that large portions of human cognitive processing are excluded. Any cognitive activity that takes place “off-line,” in the absence of task relevant input and output, is by definition not situated. Examples include planning, remembering, and day-dreaming, in contexts not directly relevant to the content of plans, memories, or day-dreams.

This observation is not new (e.g. Clark & Grush, 1999; Grush, 1997), but given the rhetoric currently to be found in the situated cognition literature, the point is worth emphasizing. By definition, situated cognition involves interaction with those things that the cognitive activity is about. Yet one of the hallmarks of human cognition is that it can take place decoupled from any immediate interaction with the environment. We can lay plans for the future, and think over what has happened in the past. We can entertain counterfactuals to consider what might have happened if circumstances had been different. We can construct mental representations of situations we have never experienced, based purely on linguistic input from others. In short, our ability to form mental representations about things that are remote in time and space, which is arguably the *sine qua non* of human thought, in principle cannot yield to a situated cognition analysis.

An argument might be made, though, that situated cognition is nevertheless the bedrock of human cognition due to our evolutionary history. Indeed, it is popular to try to drive intuitions about situated cognition by invoking a picture of our ancestors relying almost entirely on situated skills. Before we got civilized, the argument goes, the survival value of our mental abilities depended on whether they helped us to act in direct response to immediate situations such as obtaining food from the environment or avoiding predators. Thus, situated cognition may represent our fundamental cognitive architecture, even if this is not always reflected in the artificial activities of our modern world.

This view of early humans, though, most likely exaggerates the role of these survival-related on-line activities in the daily lives of early humans. With respect to obtaining food, meat-eating was a late addition to the human repertoire, and even after the onset of hunting, the large majority of calories were probably still obtained from gathering. Evidence for this claim comes from both the fossil record and the dietary patterns of hunter/gatherers today (Leaky, 1994), as well as from the dietary patterns of our nearest relatives, the chimpanzees and bonobos (de Waal, 2001). It might be more appropriate, then, to consider gathering when trying to construct a picture of our cognitive past. But gathering lends itself much less

well to a picture of human cognition as situated cognition. Successful gathering might be expected to benefit a great deal from human skills of reflective thought – remembering the terrain, coordinating with one's fellow gatherers, considering the probable impact of last week's rain, and so on. During the actual act of gathering, though, it is not clear what situated cognitive skills humans would bring to bear beyond those possessed by any foraging animal. (Put in this light, we can see that even hunting, early human style, probably involved considerable non-situated mental activity as well.)

In addition to chasing food, though, being chased by predators is also supposed to have been a major shaping force, according to this picture of the early human as a situated cognizer. Yet while avoiding predators obviously has a great deal of survival value, the situated skills of fight-or-flight are surely ancient, shared with many other species. Again, it is not clear how much mileage can be gotten out of trying to explain human intelligence in these terms. Instead, the cognitive abilities that contributed to uniquely human strategies for avoiding predation were probably of quite a different sort. As early humans became increasingly sophisticated in their social abilities, avoiding predation almost certainly involved increasing use of off-line preventative and communicative measures.

Finally, we should consider the mental activities that are known to have characterized the emerging human population and that set them apart from earlier hominid species. These included increasing sophistication of tool-making, particularly the shaping of tools to match a mental template; language, allowing communication about hypotheticals, past events, and other non-immediate situations; and depictive art, showing the ability to mentally represent what is not present, and to engage in representation for representation's sake rather than for any situated functionality (see Leakey, 1994, for further details). All of these abilities reflect the increasingly off-line nature of early human thought. To focus on situated cognition as the fundamental principle of our cognitive architecture is thus to neglect these species-defining features of human cognition.

A few counter-arguments to this can be found in the literature. Barsalou (1999a), for example, suggests that language was used by early humans primarily for immediate, situated, indexical purposes. These situated uses of language were intended to influence the behavior of others during activities such as hunting, gathering, and simple manufacturing. However, some of the examples Barsalou gives of situated uses of language appear to be in fact off-line uses, where the referent is distant in time or space – for example, describing distant terrain to people who have never seen it. One can easily think of further non-situated uses of language which would serve adaptive functions for early humans: absorbing parental edicts about avoiding dangerous behaviors; holding in mind instructions for what materials to go fetch when helping with tool manufacturing; deciding whether to join in a planned activity such as going to the river to cool off; and comprehending gossip about members of the social hierarchy who are not present. It seems plausible, then, that language served off-line functions from early on. Indeed, once the representational capacity of language emerged, it is unclear why its full capacity in this respect would not be used.

Along different lines, Brooks (1991a, p. 81) argues that, because non-situated cognitive abilities emerged late in the history of animal life on this planet, after extremely long periods in which no such innovations appeared, these were therefore the easy problems for evolution to solve (and hence, by implication, not of much theoretical interest). In fact, exactly the opposite can be inferred. Easy evolutionary solutions tend to arise again and again, a process known as convergent evolution. In contrast, the late emergence and solitary status of an animal with abilities such as manufacturing to a mental template, language, and artistic depiction, attests to a radical and complex innovation in evolutionary engineering.

In short, an argument for the centrality of situated cognition based on the demands of human survival in the wild is not strongly persuasive. Further, overstating the case for situated cognition may ultimately impede our understanding of those aspects of cognition that are in fact situated. As will be discussed in the next two sections, there is much to be learned about the ways that we engage in cognitive activity that is tightly connected with our ongoing interaction with the environment. Spatial cognition, in particular, tends to be situated. Trying to fit a piece into a jigsaw puzzle, for example, may owe more to continuous re-evaluating of spatial relationships that are being continuously manipulated than it does to any kind of disembodied pattern matching (cf. Kirsh and Maglio, 1994). For certain kinds of tasks, in fact, humans may actively choose to situate themselves (see Section 3).

Claim 2: Cognition is Time-Pressured

The previous section considered situation cognition simply to mean cognition that is situation-bound. There appears to be more, though, that is often meant by “situated cognition.” It is frequently stated that situated agents must deal with the constraints of “real time” or “runtime” (e.g. Brooks, 1991b; Pfeifer & Scheier, 1999, Chapter 3; van Gelder & Port, 1995). These phrases are used to highlight a weakness of traditional artificial intelligence models, which are generally allowed to build up and manipulate internal representations of a situation at their leisure. A real creature in a real environment, it is pointed out, has no such leisure. It must cope with predators, prey, stationary objects, and terrain, as fast as the situation dishes them out. The observation that situated cognition takes place “in real time,” is, at bottom, an observation that situated cognition must cope with time pressure.

A belief in the importance of time pressure as a shaping force in cognitive architecture underlies much of the situated cognition literature. For example, in the field of behavior-based robotics, “autonomous agents” have been built to perform tasks such as walking on an uneven surface with six legs (Quinn & Espenschied, 1993), brachiating or swinging “branch to branch” like an ape (Saito & Fukuda, 1994), and navigating around a cluttered environment looking for soda cans without bumping into anything (Mataric, 1991). Each of these activities requires real-time responsiveness to feedback from the environment. And while these activities are not especially “intelligent” in and of themselves, it is claimed that greater cognitive complexity can be built up from successive layers of procedures for real time interaction with the environment (for reviews see Brooks, 1999; Clark, 1997; Pfeifer & Scheier, 1999).

A similar emphasis on time pressure as a principle that shapes cognition can be seen as well in human behavioral research on situated cognition. For example, Kirsh and Maglio (1994) have studied the procedures people use in making time-pressured spatial decisions while playing the video game Tetris (discussed in more detail in Section 3). This research is conducted with the assumption that situations like Tetris-playing are a microcosm that can elucidate general principles of human cognition.

One reason that time pressure is thought to matter is that it creates what has been called a “representational bottleneck.” When situations demand fast and continuously evolving responses, there may simply not be time to build up a full-blown mental model of the environment, from which to derive a plan of action. Instead, it is argued, being a situated cognizer requires the use of cheap and efficient tricks for generating situation-appropriate action on the fly. (In fact, a debate has raged over whether a situated cognizer would make use of internal representations at all – Agre, 1993; Beer, 2000; Brooks, 1991a; Markman & Dietrich, 2000; Vera & Simon, 1993a,b.) Thus, taking real-time situated action as the starting point for cognitive activity is argued to have far-reaching consequences for cognitive architecture.

The force of this argument, though, depends upon the assumption that actual cognizers (humans, for example) are indeed engineered so as to circumvent this representational bottleneck, and are capable of functioning well and “normally” in time-pressured situations. But while one might wish an ideal cognitive system to have solved the problem, the assumption that *we* have solved it is disputable. Confronted with novel cognitive or perceptuo-motor problems, humans predictably fall apart under time pressure. That is, we very often do *not* successfully cope with the representational bottleneck. Lift the demands of time pressure, though, and some of the true power of human cognition becomes evident. Given the opportunity, we often behave in a decidedly off-line way: stepping back, observing, assessing, planning, and only then taking action. It is far from clear, then, that the human cognitive system has evolved an effective engineering solution for the real-time constraints of the representational bottleneck.

Furthermore, many of the activities we engage in in daily life, even many that are clearly situated, do not inherently involve time pressure. Cases include mundane activities, such as making sandwiches and paying bills, and more demanding cognitive tasks, such as doing crossword puzzles and reading scientific papers. In each of these cases, input from and output to the environment is necessary, but is at the leisure of the cognizer. (Of course, any task can be performed in a hurry, and many often are. But the state of “being in a hurry” is one that is cognitively self-imposed, and such tasks are generally performed only as fast as they can be, even if it means being late.) Situations where time-pressure is inherently part of the task, such as playing video-games or changing lanes in heavy traffic, may actually be the exception.

This is not to say, though, that an understanding of real-time interaction with the environment has nothing to contribute to our understanding of human cognition. A number of important domains may indeed be illuminated by considering them from this standpoint. The most obvious of these is perceptual-motor coordination of any kind. Even such basic activities as walking require continuous

reciprocal influence between perceptual flow and motor commands. Skilled hand movement, particularly the manipulation of objects in the environment, is another persuasive example of a time-locked perceptuo-motor activity. Building on this, more sophisticated forms of real-time situated cognition can be seen in any activity that involves continuous updating of plans in response to rapidly changing conditions. Such changing conditions often involve the activity of another human or animal that must be reckoned with. Examples include playing a sport, driving in traffic, and roughhousing with a dog. As interesting as the principles governing these cases may be in their own right, though, the argument that they can be scaled up to provide the governing principles of human cognition in general appears to be unpersuasive.

Claim 3: We Off-Load Cognitive Work onto the Environment

Despite the fact that we frequently choose to run our cognitive processes off-line, it is still true that in some situations we are forced to function on-line. In those situations, what do we do about our cognitive limitations? One response, as we have seen, is to fall apart. However, humans are not entirely helpless when confronting the representational bottleneck, and two types of strategies appear to be available when confronting on-line task demands. The first is to rely on pre-loaded representations acquired through prior learning (discussed further in Section 6). What about novel stimuli and tasks, though? In these cases there is a second option, which is to reduce the cognitive workload by making use of the environment itself in strategic ways – leaving information out there in the world to be accessed as needed, rather than taking time to fully encode it; and using *epistemic actions* (Kirsh & Maglio, 1994) to alter the environment in order to reduce the cognitive work remaining to be done.

(The environment can also be used as a long-term archive, such as in the use of reference books, appointment calendars, and computer files. This can be thought of as off-loading to avoid memorizing, which is subtly but importantly different from off-loading to avoid encoding or holding active in short term memory what is present in the immediate environment. It is the latter case that is usually discussed in the literature on off-loading. While the archival case certainly constitutes off-loading, it appears to be of less theoretical interest. The observation that we use such a strategy does not seem to challenge or shed light on existing theories of cognition. The present discussion will therefore be restricted to what we may call the situated examples of off-loading, which are the focus of the literature.)

Some investigators have begun to examine how off-loading work onto the environment may be used as a cognitive strategy. Kirsh and Maglio, as noted earlier, report a study involving the game Tetris, in which falling block shapes must be rotated and horizontally translated to fit as compactly as possible with the shapes that have already fallen. The decision of how to orient and place each block must be made before the block falls too far to allow the necessary movements. The data suggest that players use actual rotation and translation movements to simplify the problem to be solved, rather than mentally computing a solution and then executing it. A second example comes from Ballard, Hahoe, Pook and Rao (1997), who asked subjects to reproduce patterns of colored blocks under time pressure by dragging randomly scattered blocks on a computer screen into a work area and arranging them there. Recorded eye movements showed repeated referencing of the blocks in the model pattern, and these eye movements occurred at strategic moments, for example to gather information first about a block's color and then later about its precise location within the pattern. The authors argue that this is a "minimal memory strategy," and they show that it is the strategy most commonly used by subjects.

A few moments thought can yield similar examples from daily life. Not all of them involve time pressure, but other cognitive limitations, such as those of attention and working memory, can drive us to a similar kind of off-loading strategy. One example, used earlier, is physically moving around a room to generate solutions for where to put the furniture. Other examples include laying out the pieces of something that requires assembly in roughly the order and spatial relationships they will have in the finished product, or giving directions for how to get somewhere by first turning one's self and one's listener in the appropriate direction. Glenberg and Robertson (1999) have experimentally studied one such example, showing that in a compass-and-map task, subjects who were allowed to indexically link written instructions to objects in the environment during a learning phase performed better during a test phase than subjects who were not, both on comprehension of new written instructions and on performance of the actual task.

As noted earlier, this kind of strategy seems to apply most usefully to spatial tasks in particular. But is off-loading strictly limited to cases where we manipulate spatial information? Spatial tasks are only one arena of human thought. If off-loading is useful only for tasks that are themselves spatial in nature, then its range of applicability as a cognitive strategy is limited.

In fact, though, potential uses of off-loading may be far broader than this. Consider, for example, such activities as counting on one's fingers, drawing Venn diagrams, and doing math with pencil and paper. Many of these activities are both situated and spatial, in the sense that they involve the manipulation of spatial relationships among elements in the environment. The advantage is that by doing actual, physical manipulation, rather than computing a solution in our heads, we save cognitive work. However, unlike the previous examples, there is also a sense in which these activities are not situated. They are performed in the service of cognitive activity about something else, something not present in the immediate environment.

Typically, the literature on off-loading has focused on cases where the world is being used as "its own best model" (Brooks, 1991a, p. 139). Rather than attempting to mentally store and manipulate all the relevant details about a situation, we physically store and manipulate those details out in the world, in the very situation itself. In the Tetris case, for example, the elements being manipulated do not serve as tokens for anything but themselves, and their manipulation does not so much yield information about a solution as produce the goal state itself through trial and error. In contrast, actions like diagramming represent quite a different sort of use of the environment. Here, the cognitive system is exploiting external resources to achieve a solution or a piece of knowledge whose actual application will occur at some later time and place, if at all.

Notice what this buys us. This form of off-loading – what we might call *symbolic off-loading* – may in fact be applied to spatial tasks, as in the case of arranging tokens for armies on a map; but it may also be applied to non-spatial tasks, as in the case of using Venn diagrams to determine logical relations among categories. When the purpose of the activity is no longer directly linked to the situation, it also need not be directly linked to spatial problems – physical tokens, and even their spatial relationships, can be used to represent abstract, non-spatial domains of thought. The history of mathematics attests to the power behind this decoupling strategy. It should be noted, too, that symbolic off-loading need not be deliberate and formalized, but can be seen in such universal and automatic behaviors as gesturing while speaking. It has been found that gesturing is not epiphenominal, nor even strictly communicative, but seems to serve a cognitive function for the speaker, helping to grease the wheels of the thought process that the speaker is trying to express (e.g. Iverson & Goldin-Meadow, 1998; Krauss, 1998). As we will see in Section 6, the use of bodily resources for cognitive purposes not directly linked to the situation has potentially far reaching consequences for our understanding of cognition in general.

Claim 4: The Environment is Part of the Cognitive System

The insight that the body and the environment play a role in assisting cognitive activity has led some authors to assert a stronger claim: that cognition is not an activity of the mind alone, but instead is distributed across the entire interacting situation, including mind, body, and environment (e.g. Beer, 1995, pp. 182-183; Greeno & Moore, p. 49; Thelen & Smith, 1994, p. 17; Wertsch, 1998, p. 518; see Clark, 1998, pp. 513-516 for discussion). In fact, relatively few theorists appear to hold consistently to this position in its strong form. Nevertheless, an attraction to something like this claim permeates the literatures on embodied and situated cognition. It is therefore worth bringing the core idea into focus and considering it in some detail.

The claim is this: The forces that drive cognitive activity do not reside solely inside the head of the individual, but instead are distributed across the individual and the situation as they interact. Therefore, to understand cognition we must study the situation and the situated cognizer together as a single, unified system.

The first part of this claim is trivially true. Causes of behavior (and also causes of covert cognitive events such as thoughts) are surely distributed across the mind plus environment. More problematic is the reasoning connecting the first part of the claim with the second part. The fact that causal control is distributed across the situation is not sufficient justification for the claim that we must study a distributed system. The reason is that science is not ultimately about explaining the causality of any particular event. Instead, it is about understanding fundamental principles of organization and function.

Consider, for example, the goal of understanding hydrogen. Before 1900 hydrogen had been observed by scientists in a large number contexts, and a fair amount was known about its behavior when it interacted with other chemicals. But none of this behavior was really understood until the discovery in the 20th century of the structure of the atom, including the protons, neutrons and electrons that are its components and the discrete orbits that electrons inhabit. Once this was known, not only did all the previous observations of hydrogen make sense, but the behavior of hydrogen could be predicted in interactions with elements never yet observed. The causes of the behavior of hydrogen are always a combination of the nature of hydrogen plus the specifics of its surrounding context; yet explanatory satisfaction came from understanding the workings of the narrowly defined system that is the hydrogen atom. To have insisted that we focus on the study of contextualized behavior would probably not have led to a theoretical understanding with anything like this kind of explanatory force.

Distributed causality, then, is not sufficient to drive an argument for distributed cognition. Instead, we must ask what kind of a system it is that we are interested in studying. To answer this question, we must consider the meaning of the word “system,” as it is being used here. For this purpose, the contributions of systems theorists will be of help. (For a lucid summary of the issues discussed below, see Juarrero, 1999, Chapter 7.)

For a set of things to be considered a system in the formal sense, they must be not merely an *aggregate*, a collection of elements that stand in some relation to one another (spatial, temporal, or any other relation). These elements must in addition have properties that are affected by their participation in the system. Thus, the various parts of an automobile can be considered as a system because the action of the spark plugs affects the behavior of the pistons, the pistons affect the drive shaft, and so on.

But must all things that have an impact on the elements of a system themselves be considered to be part of the system? No. Many systems are *open* systems, existing within the context of an environment which can affect and be affected by the system. (No system short of the entire universe is truly closed, although some can be considered closed for practical purposes.) Thus, for example, an ecological region on earth can be considered a system in that the organisms in that region are integrally dependent on one another; but the sun need not be considered part of the system, nor the rivers that flow in from elsewhere, even though their input is vital to the ecological system. Instead, the ecological system can be considered to be an open system, receiving input from something outside itself. The fact that open systems are open is not generally considered problematic for their analysis, even when mutual influence with external forces is continuous.

From this description, though, it should be clear that how one defines the boundaries of a system is partly a matter of judgement, and depends upon the particular purposes of one’s analysis. Thus, the sun may not be part of the system when one considers the earth in biological terms, but it is most definitely part of the system when one considers the earth in terms of planetary movement. The issue, for any given scientific enterprise, is how best to carve nature at its joints.

Where does this leave us with respect to defining a cognitive system? Is it most natural, most scientifically productive, to consider the system to be the mind; or the mind, the body, and certain relevant elements in the immediate physical environment, all taken together? To help answer this question, it will be useful to introduce a few additional concepts regarding systems and how they function.

First, a system is defined by its *organization*, that is, the functional relations among its elements. These relations cannot be changed without changing the identity of the system. Next, systems can be described as either *facultative* or *obligate*. Facultative systems are temporary, organized for a particular occasion and disbanded readily. Obligate systems, on the other hand, are more or less permanent, at least relative to the lifetime of their parts.

We are now in a position to make a few observations about a “cognitive system” that is distributed across the situation. The organization of such a system – the functional relations among its elements, and indeed the constitutive elements themselves – would change every time the person moves to a new location or begins interacting with a different set of objects. That is, the system would retain its identity only so long as the situation and the person’s task-orientation toward that situation did not change. Such a system would clearly be a facultative system, and facultative systems like this would arise and disband rapidly and continuously during the daily life of the individual person. The distributed view of cognition thus trades off the obligate nature of the system in order to buy a system that is more or less closed.

If, on the other hand, we restrict the system to include only the cognitive architecture of the individual

mind or brain, then we are dealing with a single, persisting, obligate system. The various components of the system's organization – perceptual mechanisms, attentional filters, working memory stores, and so on – retain their functional roles within that system across time. The system is undeniably open with respect to its environment, continuously receiving input that affects the system's functioning and producing output that has consequences for the environment's further impact on the system itself. But, as in the case of hydrogen, or an ecosystem, this characteristic of openness does not compromise the system's status as a system. Based on this analysis, it seems clear that a strong view of distributed cognition – that a cognitive system cannot in principle be taken to comprise only an individual mind – will not hold up.

Of course we can reject this strong version of distributed cognition and still accept a weaker version, in which studying the mind-plus-situation is considered to be a promising supplementary avenue of investigation, in addition to studying the mind per se. Two points should be noted, though. First, taken in this spirit, the idea of distributed cognition loses much of its radical cachet. This view does not seek to revolutionize the field of cognitive science, but simply adds to the list of phenomena that the field studies. This can be seen as analogous to the way that chaos theory did not revolutionize or overturn our understanding of physics, but simply provided an additional tool that helped to broaden the range of phenomena that physics could successfully characterize. (Indeed, some examples of research on distributed topics appear to stretch the bounds of what we would recognize as cognition at all. The study of the organized behavior of groups is one such example – e.g. Hutchins, 1995.)

Second, it remains to be seen whether, in the long run, a distributed approach can provide deep and satisfying insights into the nature of cognition. If we recall that the goal of science is to find underlying principles and regularities, rather than to explain specific events, then the facultative nature of distributed cognition becomes a problem. Whether this problem can be overcome to yield theoretical insights with explanatory power is an issue that awaits proof.

Claim 5: Cognition is for Action

More broadly than the stringent criteria for situated cognition, the embodied cognition approach leads us to consider cognitive mechanisms in terms of their function in serving adaptive activity (e.g. Franklin, 1995, Ch. 16). The claim that cognition is for action has gained momentum from work in perception and memory in particular. “Vision,” according to Churchland, Ramachandran and Sejnowski, “has its evolutionary rationale rooted in improved motor control” (1994, p. 25; see also Ballard, 1996; O'Regan, 1992; Pessoa, Thompson & Noë, 1998). “Memory,” as Glenberg similarly argues, “evolved in service of perception and action in a three-dimensional environment” (1997, p. 1).

First, let's consider the case of visual perception. The traditional assumption has been that the purpose of the visual system is to build up an internal representation of the perceived world. What is to be done with this representation is then the job of “higher” cognitive areas. In keeping with this approach, the ventral and dorsal visual pathways in the brain have been thought of as the “what” and “where” pathways, generating representations of object structure and spatial relationships, respectively. In the past decade, though, it has been argued that the dorsal stream is more properly thought of as a “how” pathway. The proposed function of this pathway is to serve visually guided actions such as reaching and grasping (for review see Jeannerod, 1997; Goodale & Milner, 1992).

In support of this, it has been found that certain kinds of visual input can actually prime motor activity. For example, seeing a rectangle of a particular orientation facilitates performance on a subsequent grasping task, provided the object to be grasped shares that orientation (Craighero, Fadiga, Umiltà & Rizzolatti, 1996). This priming occurs even when the orientation of the rectangle does not reliably predict the orientation of the object to be grasped. A striking correlary is that visual input can activate covert motor representations in absence of any task demands. Certain motor neurons in monkeys that are involved in *controlling* tool use also respond to *seen* tools without any motor response on the part of the subject (Grafton, Fadiga, Arbib, & Rizzolatti, 1997; Murata, Fadiga, Fogassi, Gallese, Raso, & Rizzolatti, 1997). Behavioral data reported by Tucker and Ellis (1998) tell a similar story. When subjects indicate whether common objects (e.g. teapot, frying pan) are upright or inverted, response times are fastest when the response hand is the same as the hand that would be used to grasp the depicted object (e.g. left hand if the teapot's handle is on the left).

A similar proposal has been advanced for the nature of memory storage. Glenberg argues that the traditional approach to memory as “for memorizing” needs to be replaced by a view of memory as “the

encoding of patterns of possible physical interaction with a three-dimensional world” (1997, p. 1). Glenberg seeks to explain a variety of memory phenomena in terms of such perceptuo-motor patterns. Short term memory, for example, is seen not as a distinct memory “system” but as the deployment of particular action skills such as those involved in verbal rehearsal. Semantic memory and the formation of concepts are similarly explained in terms of embodied memory patterns, differing from episodic memory only in frequency of the pattern’s use across many situations.

This approach to memory helps to make sense of a variety of observations, formal and informal, that we conceptualize objects and situations in terms of their functional relevance to us, rather than neutrally or “as they really are.” These observations range from laboratory experiments on encoding specificity and functional fixedness, to the quip attributed to Mazlow that when all you have is a hammer everything looks like a nail, to the fanciful *umwelt* drawings of Von Uexkull (1934; reprints can be found in Clark, 1997) showing what the environment might look like to creatures with different cognitive agendas. Our understanding of the “how” system of vision suggests how this type of embodied memory might work. As we have seen from the work on priming of motor activity, the visual system can engage motor functions without resulting in immediate overt action. This is precisely the kind of mechanism that would be needed to create the perceptuo-motor patterning that Glenberg argues is the contents of memory.

The question we must ask, though, is how far this view of perception, memory, and cognition in general can take us. Can we dispense entirely with representation for representation’s sake, neutral with respect to a specific purpose or action? We needn’t look far for evidence suggesting that we cannot. To begin with, although the “how” system of perceptual processing appears to be for action, the very existence of the “what” system suggests that not all information encoding works this way. The ventral stream of visual processing does not appear to have the same kinds of direct links to the motor system that the dorsal stream does. Instead, the ventral stream goes about identifying patterns and objects, apparently engaging in perception for perception’s sake. This point is driven home if we consider some of the things this system is asked to encode. First, there are visual events, such as sunsets, that are always perceived at a distance and do not offer any opportunity for physical interaction (cf. Slater, 1997). Second, there are objects whose recognition depends upon holistic visual appearance, rather than on aspects of physical structure that offer opportunities for perceptuo-motor interaction. Human faces are the showcase example here, although the same point can be made for recognizing individuals of other categories, such as dogs or houses. Third, there is the case of reading, where sheer visual pattern recognition is paramount and opportunities for physical interaction with those patterns are virtually nil. Thus, perceptual encoding cannot be accounted for entirely in terms of direct perception-for-action processing channels.

The problems get worse when we look beyond perceptual processing to some of the broader functions of memory. Mental concepts, for example, do not always or even usually follow physical concrete properties that lend themselves to action, but instead often involve intangible properties based on folk-scientific theories or knowledge of causal history (e.g. Keil, 1989; Putnam, 1970; Rips, 1989). A classic example is that a mutilated dollar bill is still a dollar bill, but a counterfeit dollar bill is not. Similarly, cheddar cheese is understood to be a dairy product, but soy milk, which more closely resembles milk in its perceptual qualities and action affordances, is not.

In an ultimate sense, it must be true that cognition is for action. Adaptive behavior that promotes survival clearly must have driven the evolution of our cognitive architecture. The question, though, is: in what way or ways does our cognitive architecture subserve action? The answer being critiqued here is that the connections to action are quite direct: individual percepts, concepts, and memories are “for” (or are based on) particular action patterns. The evidence discussed above, though, suggests that this is unlikely to hold true across the board. An alternative view is that cognition often subserves action via a more indirect, flexible, and sophisticated strategy, in which information about the nature of the external world is stored for future use without strong commitments on what that future use might be.

In support of this, we can note that our mental concepts often contain rich information about the properties of objects, information that can be drawn on for a variety of uses that almost certainly were not originally encoded for. We are in fact capable of breaking out of functional fixedness, and do so on a regular basis. Thus, I can notice a piano in an unfamiliar room, and being a non-musician I might think of it only as having a bench I can sit on, and flat surfaces I can set my drink on. But I can also later call up my knowledge of the piano in a variety of unforeseen circumstances: if I need to make a loud noise to get everyone’s attention; if the door needs to be barricaded against intruders; or if we are caught in a blizzard

without power and need to smash up some furniture for fuel. Notice that these novel uses can be derived from a stored representation of the piano. They need not be triggered by direct observation of the piano and its affordances while entertaining a new action-based goal.

It is true that our mental representations are often sketchy and incomplete, particularly for things we have encountered only once and briefly. The literature on change-blindness, which shows that people can entirely miss major changes to a scene across very brief time lags, makes this point forcefully (see Simons & Levin, 1997, for review). But the fact that we are limited in how much we can attend to and absorb in a single brief encounter does not alter the fact that we can and do build up robust detailed representations with repeated exposure. Further, it is unclear that the sketchiness of a representation would prevent it from being a “representation for representation’s sake.” Our mental representations, whether novel and sketchy or familiar and detailed, appear to be to a large extent purpose-neutral, or at least to contain information beyond that needed for the originally conceived purpose. And this is arguably an adaptive cognitive strategy. A creature that encodes the world using more or less veridical mental models has an enormous advantage in problem-solving flexibility over a creature that encodes purely in terms of presently foreseeable activities.

Claim 6: Off-Line Cognition is Body-Based

Let’s return now to the kinds of externalized cognitive activities described in Section 3, in which we manipulate the environment to help us think about a problem. Consider the example of counting on one’s fingers. In it’s fullest form, this can be a set of crisp and large movements, unambiguously setting forth the different fingers as counters. But it can also be done more subtly, differentiating the positions of the fingers only enough to allow the owner of the fingers to keep track. To the observer, this might look like mere twitching. Imagine, then, if we push the activity inward still further, allowing only the priming of motor programs but no overt movement. If this kind of mental activity can be successfully employed to assist a task like counting, a whole new vista of cognitive strategies opens up.

Many centralized, allegedly abstract cognitive activities may in fact make use of sensorimotor functions in exactly this kind of covert way. Mental structures that originally evolved for perception or action appear to be co-opted and run “off-line,” decoupled from the physical inputs and outputs that were their original purpose, to assist in thinking and knowing. (Several authors have proposed mechanisms by which this decoupling might take place – Glenberg, 1997; Grush, 1996, 1998; Dennett, 1995, Chapter 13; Stein, 1994) In general, the function of these sensorimotor resources is to run a simulation of some aspect of the physical world, as a means of representing information or drawing inferences.

While this off-line aspect of embodied cognition has generated less attention than situated cognition, evidence in its favor has been quietly mounting for many years. Sensorimotor simulations of external situations are in fact widely implicated in human cognition.

Mental Imagery. Imagery, including not only the well studied case of visual imagery but also auditory imagery (Reisberg, 1992) and kinaesthetic imagery (Parsons et al, 1995), is an obvious example of mentally simulating external events. It is a commentary on the historical strength of the non-embodied viewpoint, then, that during the 1980’s the study of imagery was dominated by a debate over whether images were in fact image-like in any meaningful sense. An elaborate defense had to be mounted to show that imagery involves analog representations that functionally preserve spatial and other properties of the external world, rather than consisting of bundles of propositions (see; Kosslyn, 1994, for review). Today, this issue has been firmly resolved in favor of the analog nature of images, and evidence continues to mount for a close connection between imagery, which takes place in the absence of relevant external stimulation, and the machinery of ordinary perception (e.g. Farah, 1995; Kosslyn, Pascual-Leone, Felician & Campasano, 1999).

Working Memory. A second example of simulating physical events through the off-line use of sensorimotor resources is short term memory. Early models referred abstractly to “items” maintained temporarily in memory. Baddeley and Hitch (1974; Baddeley, 1986), however, built a persuasive case for a multi-component working memory system that had separate storage components for verbal and for visuospatial information, each of which was coded and maintained in something resembling its surface form. The particulars of the Baddeley model have been challenged on a variety of grounds, but, as I have argued elsewhere, some version of a sensorimotor model appears to be the only viable way to account for the large body of data on working memory (Wilson, 2001b). Early evidence for the sensorimotor nature

of working memory included effects of phonological similarity (worse memory for words that sound alike), word length (worse memory for long words), and articulatory suppression (worse memory when the relevant articulatory muscles are kept busy with another activity such as repeating a nonsense word). More recently, a similar set of effects, but in a different sensorimotor modality, have been found for working memory for sign language in deaf subjects – performance drops when to-be-remembered signs have similar handshapes or are temporally long, or when subjects are required to perform a repetitive movement with their hands (Wilson & Emmorey, 1997, 1998). Further, research on patient populations and brain imaging of normals indicates the involvement of speech perception and speech production areas of the brain in working memory rehearsal (see Wilson, 2001b, for review). Thus, working memory appears to be an example of a kind of symbolic off-loading, similar in spirit to that discussed in Section 3. However, instead of off-loading all the way out into the environment, working memory offloads information onto perceptual and motor control systems in the brain.

Episodic Memory. Long term memory, too, is tied in certain ways to our bodies' experiences with the world. The point is most obvious in the case of episodic memory. Whether or not one posits a separate episodic memory system, episodic memories are a class of memories defined by their content—they consist of records of spatio-temporally localized events, as experienced by the rememberer. Phenomenologically, recalling an episodic memory has a quality of “reliving,” with all the attendant visual, kinaesthetic, and spatial impressions. This is especially true when memories are fresh, before they have become crystalized by re-telling into something more resembling semantic memories.

Implicit Memory. Implicit memory also appears to be an embodied form of knowledge, consisting of a kind of perceptual and/or procedural fluency (e.g. Cohen, Eichenbaum, Deacedo, & Corkin, 1985; Johnston, Dark & Jacoby, 1985). Implicit memory is the means by which we learn skills, automatizing what was formerly effortful. Viewed in this light, implicit memory can be seen as a way of taking off-line some of the problems that confront the situated cognizer. We noted earlier that when humans are confronted with novel complex tasks under time pressure, the representational bottleneck comes into play and performance suffers. With practice, though, new skills become automatized, reducing cognitive load and circumventing the representational bottleneck. (See Epelboim, 1997, for evidence that automatizing a task reduces the need for off-loading work onto the environment.) In effect, prior experience allows whatever representations are necessary for task performance to be built up before the fact. This strategy involves exploiting predictability in the task situation being automatized – hence the fact that tasks with consist mapping between stimulus and response can be automatized, but tasks with varied mapping cannot (Schneider & Shiffrin, 1977).

Viewing automaticity as a way of tackling the representational bottleneck ahead of time can help to explain one of the apparent paradoxes of automaticity. Traditionally automatic processing has been considered the polar opposite of controlled processing (Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1997); yet highly automatized tasks appear to allow greater opportunity for fine-tuned control of action, as well as more robust and stable internal representations of the situation (cf. Uleman & Bargh, 1989). Compare, for example, a novice driver and an expert driver making a left turn, or a novice juggler and an expert juggler trying to keep three balls in the air. In each case, the degree of control over the details of the behavior is quite poor for the novice, and the phenomenological experience of the situation may be close to chaos. For the expert, in contrast, there is a sense of leisure and clarity, as well as a high degree of behavioral control. These aspects of automatic behavior become less mysterious if we consider the process of automatizing as one of building up internal representations of a situation that contains certain regularities, thus circumventing the representational bottleneck.

Reasoning and Problem-Solving. There is considerable evidence that reasoning and problem-solving make heavy use of sensorimotor simulation. Mental models, particularly spatial ones, generally improve problem-solving relative to abstract approaches. A classic example is the Buddhist monk problem: prove that a monk climbing a mountain from sunrise to sunset one day and descending the next day must be at some particular point on the path at exactly the same time on both days. The problem becomes trivial if one imagines the two days superimposed on one another. One instantly “sees” that the ascending monk and the descending monk must pass one another somewhere. Other examples of spatial models assisting reasoning and problem-solving abound in undergraduate cognitive psychology textbooks. Furthermore, recent work by Glenberg and colleagues explores how the construction of mental models may occur routinely, outside the context of formal problem-solving, in tasks such as text comprehension (Glenberg &

Robertson, 1999, 2000; Kaschak & Glenberg, 2000; see also commentaries on Glenberg & Robertson, 1999: Barsalou, 1999a; Ohlsson, 1999; Zwaan, 1999).

The domains of cognition listed above are all well established and non-controversial examples of off-line embodiment. Collectively, they suggest that there are a wide variety of ways in which sensory and motoric resources may be used for off-line cognitive activity. In accord with this, there are also a number of current areas of research exploring further ways in which off-line cognition may be embodied.

For example, the field of cognitive linguistics is re-examining linguistic processing in terms of broader principles of cognitive and sensorimotor processing. This approach, in radical contrast to the formal and abstract syntactic structures of traditional theories, posits that syntax is deeply tied to semantics (e.g. Langacker, 1986, 1991; Talmy, 2000; see Tomasello, 1998, for review). Of particular interest for the present purpose, this linkage between syntax and semantics rests in part on *image schemas* representing embodied knowledge of the physical world. These image schemas make use of perceptual principles such as attentional focus and figure/ground segregation in order to encode grammatical relations between items within the image schema.

A second example is an embodied approach to explaining mental concepts. We saw earlier that there are problems with trying to explain concepts as direct sensorimotor patterns. Nevertheless, it is possible that mental concepts may be built up out of cognitive primitives that are themselves sensorimotor in nature. Along these lines, Barsalou (1999b) has proposed that *perceptual symbol systems* are used to build up concepts out of simpler components that are symbolic and yet at the same time modal. For example, the concept of “chair,” rather than comprising abstract, arbitrary, representations of the components of a chair (“back,” “legs,” “seat”), may instead comprise modal representations of each of these components and their mutual relations, preserving analog properties of the thing being represented. While this example is quite concrete, the inclusion of *introspection* as one of the modalities helps to support the modal representation of concepts that we might think of as more abstract, such as feelings (e.g. “hungry”) and mental activities (e.g. “compare”).

A slightly different approach to abstract concepts is taken by Lakoff and Johnson and others, who argue that mental concepts are deeply metaphorical, based on a kind of second-order modeling of the physical world and relying on analogies between abstract domains and more concrete ones (e.g. Gibbs, Bogdanovich, Sykes & Barr, 1997; Lakoff & Johnson, 1980, 1999). As one example, consider the concept of “communication.” The internal structure of this concept is deeply parallel to our physical understanding of how material can be transferred from one container to another. The parallels include metaphorical movement of thoughts across space from one person’s head to another, metaphorical barriers preventing successful transfer (as when someone is being “thick-headed”), and so on. According to this view, our mental representation of communication is grounded in our knowledge of how the transfer of physical stuff works. Thus, even highly abstract mental concepts may be rooted, though in an indirect way, in sensory and motoric knowledge.

A third example is the role that motoric simulation may play in representing and understanding the behavior of conspecifics. Consider the special case of mental simulating something that is *imitable* – that can be mapped isomorphically onto one’s own body. Such stimuli in fact primarily consist of our fellow humans. There are good reasons to believe that this isomorphism provides a special foothold for robust and non-effortful modeling of the behavior of other people (see Wilson, 2001a, for review). Given that we are a highly social species, the importance of such modeling for purposes of imitating, predicting, or understanding others’ behavior is potentially quite profound.

We need not commit ourselves to all of these proposals in their present form in order to note that there is a general trend in progress. Areas of human cognition previously thought to be highly abstract now appear to be yielding to an embodied cognition approach. With such a range of arenas where mental simulation of external events may play a role, it appears that off-line embodied cognition is a widespread phenomenon in the human mind. The time may have come when we must consider these not as isolated pieces of theoretical advancement, but as reflecting a very general underlying principle of cognition.

Conclusions

Rather than continuing to treat embodied cognition as a single viewpoint, we need to treat the specific claims that have been advanced, each according to its own merits. One benefit of greater specificity is the ability to distinguish on-line aspects of embodied cognition from off-line aspects. The former include

those arenas of cognitive activity that are embedded in a task-relevant external situation, including cases that may involve time-pressure and may involve off-loading information or cognitive work onto the environment. In these cases, the mind can be seen as operating to serve the needs of a body interacting with a real-world situation. There is much to be learned about these traditionally neglected domains, but we should be cautious about claims that these principles can be scaled up to explain all of cognition.

Off-line aspects of embodied cognition, in contrast, include any cognitive activities in which sensory and motor resources are brought to bear on mental tasks whose referents are distant in time and space, or are altogether imaginary. These include symbolic off-loading, where external resources are used to assist in the mental representation and manipulation of things that are not present, as well as purely internal uses of sensorimotor representations, in the form of mental simulations. In these cases, rather than the mind operating to serve the body, we find the body (or its control systems) serving the mind. This takeover by the mind, and the concomitant ability to mentally represent what is distant in time or space, may have been one of the driving forces behind the runaway train of human intelligence that separated us from other hominids.

References

- Agre, P. E. (1993). The symbolic worldview: Reply to Vera and Simon. *Cognitive Science*, 17, 61-69.
- Baddeley, A. (1986). *Working memory*. Oxford: Oxford University Press.
- Baddeley, A., & Hitch, G. (1974). Working memory. In G. Bower (Ed.), *Recent advances in learning and motivation Vol. VIII* (pp. 647-667). Hillsdale NJ: Lawrence Erlbaum Associates.
- Ballard, D. H. (1996). On the function of visual representation. In K. Akins (Ed.), *Perception*. Oxford, U.K.: Oxford University Press.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *The Behavioral and Brain Sciences*, 20, 723-767.
- Barsalou, L. W. (1999a). Language comprehension: Archival memory or preparation for situated action? *Discourse Processes*, 28, 61-80.
- Barsalou, L. W. (1999b). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.
- Bechtel, W. (1997). Embodied connectionism. In D. M. Johnson (Ed), *The future of the cognitive revolution*, pp. 187-208. New York: Oxford University Press.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173-215.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4, 91-99.
- Brooks, R. (1999). Cambrian intelligence: The early history of the new AI. Cambridge, MA: The MIT Press.
- Brooks, R. (1991a). Intelligence without representation. *Artificial Intelligence Journal*, 47, 139-160.
- Brooks, R. (1991b). New approaches to robotics. *Science*, 253, 1227-1232.
- Brooks, R. (1990). Elephants don't play chess. In P. Maes (Ed.), *Designing autonomous agents* (p. 3-15). Cambridge, MA: MIT Press.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *Journal of Robotics and Automation*, 2, 14-23.
- Chiel, H. & Beer, R. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neurosciences*, 20, 553-557.
- Churchland, P. S., Ramachandran, V. S., & Sjenowski, T. (1994). A critique of pure vision. In C. Koch & J. Davis (Eds.), *Large scale neuronal theories of the brain*. Cambridge, MA: MIT Press.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (1998). Embodied, situated, and distributed cognition. In W. Bechtel, & G. Graham (Eds.), *A companion to cognitive sciences* (pp.506-517). Malden, MA: Blackwell Publishers Inc.
- Clark, A. & Grush, R. (1999). Towards a cognitive robotics. *Adaptive Behavior*, 7, 5-16.
- Cohen, N. J.; Eichenbaum, H., Deacedo, B. S., & Corkin, S. (1985). Different memory systems underlying acquisition of procedural and declarative knowledge. *Annals of the New York Academy of Sciences*, 444, 54-71.
- Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simon & Schuster.
- di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., & Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. *Experimental Brain Research*, 91, 176-180.
- Epelboim, J. (1997). Deictic codes, embodiment of cognition, and the real world. *Behavioral & Brain Sciences*, 20, 746.
- Farah, M. J. (1995). The neural bases of mental imagery. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences*, pp. Cambridge, MA: MIT Press.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Franklin, S. (1995). *Artificial Minds*. Cambridge, MA: MIT Press.
- Gallese V., Fadiga L., Fogassi L., & Rizzolatti G. (1996). Action recognition in the premotor cortex. *Brain*, 119, 593-609.
- Gibbs, R. W., Bogdanovich, J. M., Sykes, J. R., & Barr, D. J. (1997). Metaphor in idiom comprehension. *Journal of Memory and Language*, 37, 141-154.
- Glenberg, A. M. (1997). What memory is for. *The Behavioral and Brain Sciences*, 20, 1-55.
- Glenberg, A. M. & Robertson, D. A. (1999). Indexical understanding of instructions. *Discourse Processes*, 28, 1-26.

- Glenberg, A. M. & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. Journal of Memory and Language, 43, 379-401.
- Goodale, M.A. & Milner, A.D. (1992). Separate visual pathways for perception and action. Trends in Neurosciences, 15, 20-25.
- Goodwin, C. J. (1999). A history of modern psychology. New York: John Wiley & Sons.
- Greeno, J. G. & Moore, J. L. (1993). Situativity and symbols: Response to Vera and Simon. Cognitive Science, 17, 49-59.
- Grush, R. (1998). Perception, imagery, and the sensorimotor loop. www.pitt.edu/~grush/papers/%21papers.html. English translation of: Wahrnehmung, Vorstellung und die sensomotorische Schleife. In *Bewußtsein und Repraesentation*, F. Esken & H.-D. Heckmann (Eds). Paderborn, Germany: Verlag Ferdinand Schoeningh.
- Grush, R. (1997). Yet another design for a brain? Review of Port and van Gelder (Eds.), *Mind as Motion*. Philosophical Psychology, 10, 233-242
- Grush, R. (1996). Emulation and cognition. Unpublished doctoral dissertation, University of California, San Diego.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge MA: MIT Press.
- Iverson, J. M., & Goldin-Meadow, S. (1998). Why people gesture when they speak. Nature, 396, 228.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Cambridge, MA: Blackwell Publishers.
- Johnston, W. A., Dark, V. J., & Jacoby, L. L. (1985). Perceptual fluency and recognition judgments. Journal of Experimental Psychology: Learning, Memory & Cognition, 11, 3-11.
- Juarrero, A. (1999). Dynamics in action: Intentional behavior as a complex system. Cambridge, MA: MIT Press.
- Kaschak, M. P. & Glenberg, A. M. (2000). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. Journal of Memory and Language, 43, 508-529.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. Cognitive Science, 18, 513-549.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Kosslyn, S. M., Pascual-Leone, A., Felician, O., & Camposano, S. (1999). The role of Area 17 in visual imagery: Convergent evidence from PET and rTMS. Science, 284, 167-170.
- Krauss, R. M. (1998). Why do we gesture when we speak? Nature, 7, 54-60.
- Lakoff, G. & Johnson, M. (1999). Philosophy in the flesh: The embodied mind and its challenge to western thought. New York: Basic Books.
- Lakoff, G. & Johnson, M. (1980). Metaphors we live by. Chicago: University of Chicago Press.
- Langacker, R. (1986, 1991). Foundations of cognitive grammar (2 vols.). Stanford: Stanford University Press.
- Leakey, R. (1994). *The origin of humankind*. New York: BasicBooks.
- Markman, A. B. & Dietrich, E. (2000). In defense of representation. Cognitive Psychology, 40, 138-171.
- Mataric, M. (1991). Navigating with a rat brain: A neurobiologically inspired model for robot spatial representation. In J.-A. Meyer and S. Wilson (Eds.), *From animals to animats*. Cambridge MA: MIT Press.
- Ohlsson, S. (1999). Anchoring language in reality: Observations on reference and representation. Discourse Processes, 28, 93-105.
- O'Regan, K. (1992). Solving the "real" mysteries of visual perception: the world as an outside memory. Canadian Journal of Psychology, 46, 461-488.
- Pessoa, L., Thomson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. Behavioral and Brain Sciences, 21, 723-802.
- Port, R. F. & van Gelder, T. (1995). Mind as motion: Explorations in the dynamics of cognition. Cambridge, MA: MIT Press.
- Prinz, W. (1987). Ideo-motor action. In H. Heuer & A.F. Sanders (Eds.), Perspectives on perception and action (pp.47-76). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Parsons, L. M., Fox, P. T., Downs, J. H., Glass, T., Hirsch, T. B., Martin, C. C., Jerabek, P. A., &

- Lancaster, J. L. (1995). Use of implicit motor imagery for visual shape discrimination as revealed by PET. *Nature*, *375*, 54-58.
- Pfeifer, R. & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: The MIT Press.
- Port, R. F. & van Gelder, T. (1995). Mind as motion: Explorations in the dynamics of cognition. Cambridge, MA: MIT Press.
- Putnam, H. (1970). Is semantics possible? in H.E. Kiefer and M. K. Munitz (Eds.), *Language, belief and metaphysics*. Albany, NY: State University of New York Press.
- Quinn, R., & Espendchied, K. (1993). Control of a hexapod robot using a biologically inspired neural network. In R. Beer, R. Ritzman & T McKenna (Eds.), *Biological neural networks in invertebrate neuroethology and robotics*. Academic Press.
- Reisberg, D. (Ed.) (1992). Auditory Imagery. Hillsdale, NJ: Erlbaum.
- Rips, L. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*, pp. 21-59. Cambridge, UK: Cambridge University Press.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *cognitive Brain Research*, *3*, 131-141.
- Saito, F., & Fukuda, T. (1994). Two link robot brachiation with connectionist Q-learning. In D. Cliff (Ed.), *From animals to animats 3*. Cambridge MA: MIT Press.
- Schneider, W. & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*, 1-66.
- Shiffrin, R. M. & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*, 127-190.
- Simons, D. J. & Levin, D. T. (1997). Change blindness. Trends in Cognitive Sciences, *1*, 261-267.
- Slater, C. (1997). Conceptualizing a sunset \neq using a sunset as a discriminative stimulus. *Behavioral and Brain Sciences*, *20*, 37-38.
- Steels, L. & Brooks, R. (1995). The artificial life route to artificial intelligence: Building embodied, situated agents. Hillsdale, NJ: Erlbaum.
- Stein, L. (1994). Imagination and situated cognition. *Journal of Experimental Theoretical Artificial Intelligence*, *6*, 393-407.
- Talmy, L. (2000). Toward a cognitive semantics, Vol. I: Conceptual structuring systems. Cambridge, MA: MIT Press.
- Thelen, E. & Smith, L. B. (1994). A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.
- Tomasello, M. (1998). Cognitive linguistics. In W. Bechtel, & G. Graham. (Eds.), A companion to cognitive science (pp.477-487). Malden, MA: Blackwell Publishers Inc.
- Tucker, M. & Ellis, R. (1998). On the relations between seen objects and components of potential actions. Journal of Experimental Psychology: Human Perception and Performance, *24*, 830-846.
- Uleman, J. & Bargh, J. (Eds). (1989). *Unintended thought*. New York: Guilford.
- van Gelder, T. & Port, R. (1995). It's about time: An overview of the dynamical approach to cognition. In R. Port & T. van Gelder (Eds.), Mind as motion: Explorations in the dynamics of cognition. Cambridge, MA: MIT Press.
- Vera, A. H. & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science*, *17*, 7-48.
- Von Uexkull, J (1934). A stroll through the worlds of animals and men. In K. Lashley (Ed.), Instinctive behavior. International Universities Press.
- Wertsch, J. V. (1998). Mediated action. In W. Bechtel, & G. Graham (Eds.), A companion to cognitive science (pp.518-525). Malden, MA: Blackwell Publishers Inc.
- Wiles, J. & Dartnall, T. (1999). *Perspectives on cognitive science: Theories, experiments, and foundations*. Stamford, CT: Ablex Publishing Corporation.
- Wilson, M. (2001a). Perceiving imitable stimuli: Consequences of isomorphism between input and output. In press, Psychological Bulletin.
- Wilson, M. (2001b). The case for sensorimotor coding in working memory. Psychonomic Bulletin and Review, *8*, 44-57.
- Wilson, M. & Emmorey, K. (1997). A visuo-spatial "phonological loop" in working memory: Evidence from American Sign Language. *Memory & Cognition*, *25*, 313-320.

- Wilson, M. & Emmorey, K. (1998). A “word length effect” for sign language: Further evidence on the role of language in structuring working memory. *Memory & Cognition*, 26, 584-590.
- Zwaan, R. A. (1999). Embodied cognition, perceptual symbols, and situation models. Discourse Processes, 28, 81-88.