Port, R., & van Gelder, T. J. (1995). *Mind as motion: Dynamics, behavior, and cognition.* Cambridge, MA: MIT Press.

Rosenthal, S. B. (1983). Meaning as habit: Some systematic implications of Peirce's pragmatism. In E. Freeman (Ed.), *The relevance of Charles Peirce* (pp. 312–327). La Salle, IL: Open Court.

Rosenthal, S. B. (1987). Classical American pragmatism: Key themes and phenomenological dimensions. In R. S. Corrington, C. Hausman, & T. M. Seebohm (Eds.), *Pragmatism considers phenomenology* (pp. 37–57). Washington, DC: University Press of America.

Rizzolatti, G., Fadiga, L., Fogassi, L., & Gallese, V. (1997). The space around us. *Science, 277,* 190–191.

Shanon, B. (1993). *The representational and the presentational.* Hertfordshire, England: Harvester Wheatsheaf.

Smith, J. E. (1987). The reconception of experience in Peirce, James, and Dewey. In R. S. Corrington, C. Hausman, & T. M. Seebohm (Eds.), *Pragmatism considers phenomenology* (pp. 73–91). Washington, DC: University Press of America.

Thayer, H. S. (1973). *Meaning and action.* Indianapolis, IN: Bobbs-Merrill.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327–352.

Tversky, A., & Gati, I. (1978). Studies of similarity. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 79–98). Hillsdale, NJ: Lawrence Erlbaum Associates.

Prinz, J.J., & Barsalou, L.W. (2000). Steering a course for embodied representation. In E. Dietrich & A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines* (51-77). Cambridge, MA: MIT Press.

3

Steering a Course for Embodied Representation

Jesse J. Prinz
Washington University, St. Louis

Lawrence W. Barsalou
Emory University

Cognition is a dynamic process, continually changing over time. This truism has contributed to a growing interest in using dynamic systems theory to study and model cognitive phenomena. This development is by no means unwelcome. Already, dynamic approaches have inspired important new insights into cognition, and they are likely to inspire more. There is, however, need for caution. Some proponents of the dynamic approach have been too quick to dismiss convictions that have been widely held since the advent of cognitive science. Sometimes there is an unfortunate tendency to view a new approach as so revolutionary that important lessons of previous paradigms are prematurely discarded. The rapidly growing literature on dynamic approaches to cognition is rife with such revolutionary rhetoric. One of the most provocative claims is that the talk of internal representations should be seriously delimited or even eliminated from explanations of cognition (Freeman & Skarda, 1990; Thelen & Smith, 1994; van Gelder, 1995; see also Edelman, 1989).

In our opinion, representation has been a red herring in the dynamic systems revolution. In making strong eliminativist claims, dynamicists have probably obscured important ramifications of their position, as they appear increasingly to realize. For example, van Gelder (1997) de-emphasized the critique of representation in his most recent presentation of the dynamic approach and suggested that dynamic systems can accommodate representations in various ways. The critique of representations can be more instructively viewed as part of a larger campaign that dynamicists wage against the prevailing view in cognitive science. The central issue is no longer whether internal representations exist, but whether traditional representational approaches can account for the contextually sensitive, embodied, and inherently temporal character of cognition. Traditionally conceived, representations are context invariant, disembodied, and static. If representations are inextricably tied to these properties, dynamicists believe that representational theories of cognition are inadequate. If, on the other hand, representations are to be reconstrued as contextually sensitive, embodied, and temporal, then they should be identified with state-space trajectories, attractors, bifurcations, or some other construct of dynamic systems theory. In other words, the applicability of dynamic tools in studying cognition is alleged to either undermine or appropriate talk of internal representations. These two possibilities constitute two distinct challenges: an elimination challenge and an appropriation challenge.

We agree that cognition is contextually sensitive, embodied, and inherently temporal. We also agree that dynamic systems theory provides useful tools for describing cognitive systems. In that sense, we uphold the spirit of current dynamical systems approaches. At the same time, we think that internal representations should be neither eliminated nor appropriated by dynamic systems theory. First, we offer some general considerations in favor of representational approaches to cognition. In particular, we show that many dynamic systems lend themselves to representational interpretations, and we argue that representations contribute to the adaptive success of evolved cognitive systems. Then we take up context sensitivity, embodiment, and temporality in turn. While conceding the importance of these properties, we argue that none of them calls for the elimination of representations or for identification of representations with constructs posited by dynamic systems theory. After that, we introduce a perceptual symbols theory of mental representation, which can account for context sensitivity, embodiment, and temporality without being stated in terms of dynamic systems theory. None of these arguments is intended to show that dynamic tools should not be used in studying cognition. To the contrary, we conclude by proposing that dynamic systems theory and perceptual symbols theory are complementary. They can work in concert to describe different aspects of cognitive systems.

REPRESENTATION

A dynamic system is simply a physical system (or a physically realizable system) that changes over time. A mathematical description of a dynamic system contains a set of numbers that describe the system's states and a set of functions that specify how the system's current state evolves into

its next state. Given this generic characterization, it should be obvious that representational systems can be dynamic. Implemented programs in classical computer languages, for instance, are representational systems that change over time (see Giunti, 1995). At every moment, an implemented program is in some state, and the changes it undergoes can be plotted as a trajectory through a space of possible states. Likewise, many connectionist networks lend themselves to representational interpretations. Individual nodes, in the case of localist networks, or activation patterns across nodes, as in distributed networks, can often be described as representations. The changing patterns of activity in such networks are, nevertheless, dynamic and can be usefully characterized by using standard tools of dynamic systems theory.

The challenge posed by dynamic approaches to cognition comes in recognizing that, although implemented representational systems are all dynamic, many dynamic systems are nonrepresentational. Perhaps, dynamicists have argued, natural cognitive systems fall into that category. They think we should study cognition by using the tools of dynamic systems theory (e.g., low-dimensional differential equations) instead of by appealing to internal representations.

Nonrepresentational Dynamics Systems

Van Gelder (1995) provided one of the most perspicuous arguments against internal representations based on the dynamic character of natural cognitive systems. It rests on an analogy to another quintessential dynamic system: Watt's centrifugal governor. Van Gelder first proposed that the Watt governor achieves a complex engineering feat without representations. He then urged us to consider the possibility that cognition is more like a Watt governor than a Turing machine, a classic representational system. We focus on problems with the first stage of van Gelder's argument, namely, that the Watt governor is not representational. In this discussion, we do not defend the claim that cognition is computational. Further doubts about both stages of his argument can be found elsewhere (Bechtel, 1996; Clark, 1997b; Clark & Toribio, 1994).

As Fig. 3.1 illustrates, the Watt governor is designed to translate the fluctuating pressure from a steam piston into the uniform rotations of a flywheel by adjusting a throttle valve. It achieves this remarkable feat with a simple arrangement of parts. First, a spindle is attached to the flywheel, then two arms with balls at their ends are attached by hinges to the spindle, and finally the arms are linked to the throttle valve. As steam pressure increases, the flywheel's rotation causes the spindle to rotate faster; the spindle's rotation causes the arms to move outward by centrifugal force; and the arms' outward motion closes the throttle valve, decreasing pres-

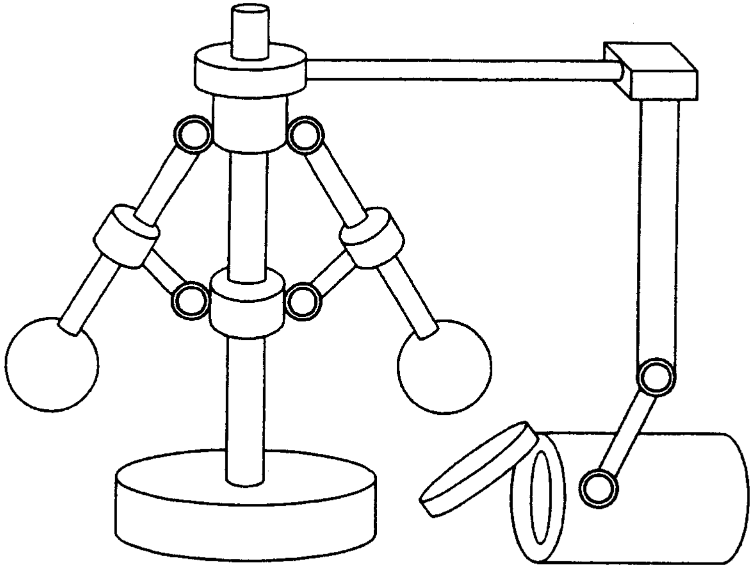

FIG. 3.1. A centrifugal governor.

sure. When steam is decreased, the flywheel slows, the arms fall, and the valve opens again, increasing pressure. By modulating pressure in this manner, the flywheel maintains a relatively constant speed.

On our definition and van Gelder's, the Watt governor is a dynamic system, a physical device that changes continuously in time. Like other dynamic systems, its behavior can be described by using differential equations, which van Gelder provided. These equations do not include variables or parameters that correspond in any obvious way to representations. However, these equations neither demonstrate nor corroborate van Gelder's claim that the Watt governor is not a representational device.

Van Gelder anticipated an obvious objection to this claim by acknowledging the "initially quite attractive intuition" (1995, p. 351) that arm angle of a centrifugal governor represents flywheel speed. Despite this intuition, van Gelder insisted that the relation between arm angle and engine speed is not representational. To support this claim, he provided four arguments that "are not based on an unduly restrictive notion of representation." Bechtel (1996) criticized these arguments effectively, so we do not review them all here. Instead, we focus on positive reasons for believing that the Watt governor harbors representation.

To make this case, we first need to determine what it means to represent. On many current accounts, representation is defined with synonyms. Something is a representation if it denotes, designates, stands for, refers to, or is about something else. Debate rages over what conditions must

be met to stand in this kind of relation (Stich & Warfield, 1994). Entering into this debate would take us too far afield. Instead, we simply describe one theory of representation that epitomizes the prevailing family of views. If the Watt governor counts as representational on this account, we have reason to doubt van Gelder's diagnosis.

Many philosophers believe that representation involves information. A state represents a property only if it carries information about it (Dretske, 1981; Fodor, 1990). Carrying information is, in turn, analyzed in terms of nomic, or law-like, covariation. To a first approximation, a state *s* carries information about a property *F* just in case instantiations of *F* reliably cause tokens of *s*. Although arguably necessary, information is not sufficient for representation. Fire reliably causes smoke, but smoke does not represent fire. An information-bearing state must satisfy one further condition to count as a representation. For Dretske (1995), this extra ingredient is teleological: Something can represent something else only if it has the *function* of carrying information about it. Representations can acquire such functions through a variety of different histories, including natural selection, learning, and design. For example, the word *fire* has the function of carrying information about fires, because it was created to linguistically indicate their presence. In contrast, smoke, even though it covaries with fires, does not have the function of carrying information about them and therefore does not represent them.

In sum, Dretske said that two conditions must be met for something to represent something else. First, the former must carry information about it through covariance. The Watt governor certainly satisfies this condition. To say that arm angles carry information about engine speed just means that arm angles causally covary with engine speed in a law-like way. Van Gelder revealed the satisfaction of this condition when he provided a law correlating changes in arm angle with engine speed (1995, p. 356). Second, to count as a representation on Dretske's account, an information-bearing state must have the function of bearing that information. Once again, the Watt governor complies. Watt *designed* the governor to regulate the relation between steam pressure and engine speed. Most important, he designed the arm angles to covary with engine speed, to release the right amount of steam. Because design confers function, arm angle *represents* engine speed. In Dretske's framework, it is the function of arm angles to carry information about engine speed.

We do not necessarily claim that Dretske's theory is correct in all its details. What is crucial is that it is exemplifies a class of accounts that currently dominate philosophical approaches to representation. Because the general form of this account is readily applicable to the states of many dynamic systems, we have good reason to believe that these systems are representational.

Van Gelder might complain that this analysis fails to address one of his arguments against representational interpretations of the Watt governor. On this particular argument, representation is the wrong conceptual tool for explaining the relation between arm angle and engine speed. Instead, this relation is better understood as an instance of *coupling*, a central construct in dynamics "more subtle and complex" than representation (1995, p. 353). Van Gelder considered this his most important argument, and it is the only one included in his most recent treatment of representation (van Gelder, 1997).

We find this argument puzzling. Why must coupling be incompatible with representation? To simply assume that it is begs the question. On examining the construct of coupling more closely, no incompatibility is apparent. To see this, consider Dretske's characterization of thermostats as paradigmatic representational systems, in particular, thermostats that use bimetal strips to control temperature (Dretske, 1988, p. 49). Each side of a bimetal strip contains a different temperature-sensitive metal. When the strip cools, the different expansion rates of the two metals cause it to bend in one direction; when the strip warms up, they cause it to bend in the opposite direction. As Fig. 3.2 illustrates, bending in these two directions turns a furnace on or off. As the temperature of the room falls, the metal strip bends toward an electrical contact, eventually touching it and closing a circuit to activate the furnace. As the furnace warms the room, the metal strip bends in the other direction, thereby opening the circuit and deactivating the furnace.

As this example illustrates, the thermostat and room temperature constitute coupled dynamic systems. The dynamic behavior of the thermostat is coupled with dynamic changes in room temperature. Most important, however, this relation is also representational. On the generic account of representation just provided, the thermostat carries information about room temperature by functional design. Thus, the relation between the thermostat and room temperature simultaneously constitutes examples of coupled systems *and* representation.

The compatibility of coupling and representation can be illustrated by another example. In cognitive neuroscience, it is widely supposed that the brain can be functionally decomposed into different systems (Mundale & Bechtel, 1996; Van Essen, Felleman, DeYoe, Olavarria, & Knierim, 1990; Zeki, 1992). These systems are often presumed to carry out representational functions. For example, different portions of the visual system are said to represent different aspects of a visual signal. Cells in brain area V3 respond to dynamic form, area V4 responds to color and color form, and area MT responds to motion. Following Dretske, these cells can be said to represent those properties in virtue of the fact that they are reliably caused by them. Such reliable causal relations are exactly what neuroscientists appeal to in mapping the cortex. At the same time, different regions of the visual system are thought to interact. Direct and indirect pathways between these regions give rise to a complex coupled system. Activity in one area can both affect and by affected by its neighbors. Such coupling can play various roles. For example, dynamic coupling can allow for temporal synchronization, which, in turn, may be necessary for generating a unified conscious percept (Crick & Koch, 1990). As in the thermostat, representation and coupling are compatible and complementary in the nervous system. Omitting either of these components from one's analysis significantly diminishes understanding. Each underwrites distinct, but equally important and often interrelated, explanatory projects.

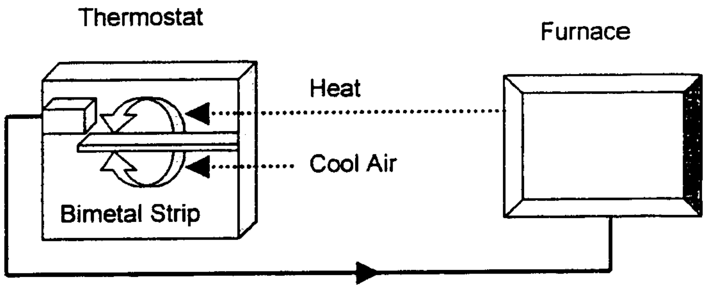

FIG. 3.2. A bimetal thermostat.

The upshot of this discussion is that many systems that readily lend themselves to dynamic analysis also lend themselves to representational analysis. It is sometimes thought that the mere applicability of dynamic tools, such as differential equations, tells against representational analyses of a system. This is not the case. Many interesting dynamic systems, such as Watt governors and thermostats, *are* representational. Furthermore, it appears that construing these dynamic systems as representational offers an important and essential level of analysis in characterizing them. If governor arms did not carry information about flywheel speed and thermostat strips did not carry information about temperature, they would not be able to carry out the functions for which they were designed. To the extent that dynamicists believe that cognition is anything like Watt governors and thermostats, it follows that cognition, too, is inherently representational.

The Evolutionary Fitness of Representation

We have argued that the dynamic systems sometimes used to characterize cognition are inherently representational. However, our final remarks alluded to a much stronger point: Representation is extremely useful. In the case of natural cognitive systems, we believe that evolution has selected representation as a central design function of cognition, because representation confers substantial advantages in fitness on organisms that implement it powerfully. In the spirit of our earlier discussion, we believe that representational ability is analogous to other high-level functions that organisms exhibit, such as reproduction and digestion. Because these functions increase fitness, evolution attempts to implement them in organisms.

In his evolutionary theory of cognition, Donald (1991, 1993) elegantly postulated an important relation between representational ability and fitness. Donald argued that all nonhuman species are relatively limited in their representational abilities. To a large extent, nonhuman organisms represent only their current environment and internal state. Although these organisms may anticipate short-term events that are not present (e.g., the future acquisition of prey on initially perceiving it), they do not typically represent nonpresent situations that have nothing to do with current circumstances.

The major transition leading to human intelligence, Donald continued, occurred when humans acquired powerful new abilities for representing nonpresent situations. By being able to review past events and preview future events, humans became better able to learn from the past and to control the future. Furthermore, as humans developed new communicative abilities for inducing representational states in others, their cooperative activities gained even further leverage on environmental resources. Together, multiple individuals could review past events, control current events, and plan future events. In addition, humans developed extensive cultural systems capable of accomplishing much more than a single individual could accomplish alone.

In our opinion, dynamic accounts of cognition that eschew representation essentially create organisms at Donald's first level of intelligence. Because these organisms have minimal representational ability, they are relatively limited in their ability to represent nonpresent situations. Instead, they simply have the reactive ability to encode the current situation and to generate appropriate actions. Certainly, these abilities are important and sophisticated. Nevertheless, an organism with only these abilities is much less powerful than organisms that can represent and manipulate nonpresent situations. Like Donald, we believe that evolution selected increasingly powerful representational abilities to create an unusually fit species, namely, humans.

Roles of Representation in Human Cognition

As just discussed, we believe that representation plays one essential role of "standing in" for situations that are not present in the physical world. For example, by jointly representing a herd of buffalo traversing a valley, a group of hunters can review past hunting episodes and plan future ones. Rather than simply reacting to an actual herd, the hunters can imagine various scenarios that may be factual or counterfactual. By considering these scenarios systematically and evaluating their properties, the hunters can converge on optimal hunting practices. Most important, because these representations correspond systematically to actual situations, they support offline simulations that are truly effective in guiding future action.

A second essential role of representation is to support productivity. In cognitive science, it is widely agreed that humans have the ability to construct infinite representations from a finite set of building blocks. Not only does productivity manifest itself in the infinite production of sentences from finite words (Chomsky, 1957), it also manifests itself in the infinite production of conceptualizations from finite concepts (Fodor & Pylyshyn, 1988). Having such an unbounded capacity is invaluable, because it allows us to cope with an environment that is constantly changing. Anticipating and understanding new situations require us to recombine previous knowledge in novel ways.

We agree with Fodor and Pylyshyn when they criticize dynamic models (qua connectionism) that fail to account for this basic property of human cognition. We agree that any viable theory of cognition must exhibit productivity. However, we strongly disagree with other aspects of Fodor and Pylyshyn's proposal. We reject their underlying commitments to classical representations that are static, context invariant, and disembodied. In other words, we embrace only Fodor and Pylyshyn's emphasis on the importance of combinatorial representations as a means of achieving productivity. Dynamic approaches that repudiate productive representations give up valuable resources for explaining our ability to cope with a changing environment. Many connectionists agree and have tried to meet Fodor and Pylyshyn's challenge by generating networks that use combinatorial representations. For reasons to be discussed below, we have reservations about many connectionist accounts. We support these connectionists, however, for conceding the importance of representations and for seeking alternatives to traditional, static accounts.

CONTEXT SENSITIVITY

Cognitive science has traditionally conceived of representations as enduring, context-insensitive entities. On this view, the representation of a category may be a single node, a stable pattern of activation, a fixed set of features, a constant set of predicates, a static network, an immutable mental word, and so forth. This approach assumes that an invariant structure represents a category in all contexts. Different contexts do not require that context-sensitive representations be tailored to fit them. This picture seems deeply problematic. Natural cognitive systems exhibit remarkable variability. Our behaviors and inner states adapt differently to different contexts. This suggests that the traditional concept of representations is inadequate. Dynamic systems theorists claim that our extreme sensitivity to context is best captured by using the tools of dynamic systems theory. Some have even taken context sensitivity as support for the mandate to banish representational talk from our explanatory repertoire (Thelen & Smith, 1994). Although we are sympathetic to the underlying spirit of these arguments, we think this case is overstated. Context sensitivity coexists with a degree of stability that is best captured by appeals to representations.

Context Sensitivity in Dynamic Systems

Complexly coupled dynamic systems often process inputs of the same type in different ways. Because such systems' behaviors are a product of activities in numerous interconnected components, they are typically incapable of identically reproducing an earlier state. Even when a focal point in the environment contains the same input as it did on a previous occasion, active inner states and peripheral environmental features make up a distinct context. These influence the way the focal input is processed. Thus, when confronted with exactly the same focal input, the system responds to it somewhat differently. In these ways, dynamic systems produce context-sensitive responses to inputs, making these systems significantly different from more traditionally conceived systems that process identical inputs in exactly the same way.

One crusader for context specificity is Walter Freeman. Using the tools of dynamic systems theory, Freeman has explored the neural processing of categories and has found that the neural states underlying categories are far more context specific than previously suspected (Freeman, 1991). On the basis of these findings, Freeman has concluded that cognition does not traffic in representations, which he assumed are inherently stable.

As an example of Freeman's important findings, consider his work on rabbit olfaction. Freeman demonstrated that rabbit olfactory states change dramatically as rabbits learn to distinguish new odorants. Using electroencephalogram recordings, Freeman showed that different odorants are associated with different spatial patterns in the amplitudes of waves across a rabbit's olfactory bulb. Repeated exposure to a given training stimulus, such as sawdust, invokes the same pattern under the same conditions. This continuity is ephemeral, however. When a rabbit is conditioned to another stimulus, such as bananas, the patterns associated with previously learned odorants change. Freeman explained this by appeal to the distributed and interconnected nature of neural states. The neural state for an odorant on a given occasion does not reflect only the stimulus; it also reflects the rabbit's learning history. If the neural state corresponding to an odor changes across learning contexts, the traditional faith in enduring, invariant representations cannot be maintained.

Freeman's results provide a fascinating look into how the brain works. A population of neurons mediating one odor engages in complex interactions with populations of neurons mediating other odors, producing dynamic changes in each. However, dynamic variability at the neural level does not froth upward into functional properties of the system. Indeed, Freeman's behavioral results clearly indicate that dynamic variability is confined to the neural level. At the behavioral level, a rabbit's response to a particular odor remains constant while the underlying neural states vary. Different patterns in a rabbit's olfactory bulb mediate the same functional relation between an odor input and a behavioral output. In other words, these different neural states serve the same input–output (I/O) function.

This observation suggests that one must be cautious when arguing for context specificity. Observing radical variability at the neuronal level can give the appearance that cognitive systems are contextually sensitive all the way up. From this perspective, it looks as if there is no level at which one can postulate stable representations, and, as a result, the very notion of a representation seems to get no foothold. If this were correct, it would be appropriate to treat the brain as a purely nonrepresentational dynamic system. The fact that neuronally distinct states serve common I/O functions points to another interpretation: Although there is extreme context sensitivity at the lowest level of analysis, there is also a somewhat higher level of analysis at which stability emerges through function. For certain explanatory purposes, it is appropriate to treat neurally distinct states as the same precisely because of their functional equivalence. If we did not co-classify such states, it would be extraordinarily difficult to explain how we manage to recognize the same inputs again and again. Thus, the applicability of dynamic tools at the lowest level does not vitiate a representational account at a somewhat higher level. Indeed, it seems to demand high-level analysis if we want an explanatory psychology.

The longstanding importance of emergent properties in science constitutes a good reason for bestowing a stable representational interpretation on the neural states that implement a common I/O function. Consider emergent properties such as reproduction, digestion, flying, and life. Just because multiple physical states instantiate reproduction, for example, does not mean that we eschew it as a useful scientific construct. It is indeed interesting and significant that reproduction can be achieved in many ways, yet it remains critical that the same state is achieved in all cases. Indeed, one can argue that evolution selects for these higher level properties and cares more about achieving and preserving them than about the particular mechanisms that implement them. As a result, different life forms achieve the same ability but in ways that vary widely. Eschewing representations because they supervene on disparate lower level activities makes as much sense as eschewing reproduction and digestion on the same grounds.

Context Sensitivity in Representational Models

To this point, we have argued that low-level context sensitivity is compatible with higher level stability. This suggests that complex neural dynamics cannot refute representational accounts of the brain. At the same time, it is important to show that representational accounts are compatible with context sensitivity. If stable representations precluded this, representational theories of cognition would be intolerably inflexible. We believe that representational accounts can readily exhibit context sensitivity. We also believe that the context sensitivity of representational accounts can be neatly interfaced with context sensitivity in dynamic systems models to provide unified multilevel accounts of this important property.

Barsalou (1987, 1989, 1993) demonstrated that concepts exhibit considerable context sensitivity. When representing the same category, different subjects represent it with different information. In the feature-listing task, for example, the features active for one subject overlap only about 40% with the features active for another subject. Similarly, when the same individual represents a category on different occasions, he or she represents it differently. In the feature-listing task, the features that a subject produces for a category on one occasion overlap only about 65% with the features produced 2 weeks later. Finally, context substantially biases the features active for a subject on a given occasion (Barsalou, 1982; Medin, Goldstone, & Gentner, 1993). For example, subjects do not usually activate the property *pet* for snake but do so in the context of other pets. All these results clearly document the context sensitivity that dynamicists champion.

All these discoveries were made and explained in a representational framework (e.g., Barsalou, 1989, in press). For example, one can view an individual's knowledge of a category as containing many representations of features and relations that have accrued over a lifetime of experience with category members. On a given occasion, only a very small subset of this representational base becomes active to represent the category. In particular, the features that become active are those that have been processed most frequently, those that have been processed most recently, and those that are most associated to the current context. Note that frequency, recency, and context are exactly the sorts of statistical factors central to dynamic theories. However, they have been conceptualized and implemented in the representational framework. Indeed, such flexibility has become a hallmark of many representational models. For example, the prototypes of prototype models and the exemplars of exemplar models are readily viewed as context-sensitive representations whose specific form varies widely across contexts (e.g., Barsalou, 1990). On this approach, context sensitivity is achieved not by eradicating stability, but by creating novel and adaptive combinations of relatively stable feature representations.

The fundamental error of more traditional representational theories is not stability as such, but too much stability. On such theories, categories are represented by the very same collection of features on each occasion. On the present view, categories are represented by different collections of features as context demands. On both accounts, features are construed as representations, and they can be relatively stable entities stored in memory. The latter account improves on the former, however, by insisting that features be retrieved, combined, and deployed in ways that vary. From this perspective, one can capture the dynamists' intuition that our cognitive systems are never in the same state twice without giving up the notion of representation.

Prinz (1997) developed one further argument reinforcing the message that context sensitivity does not vitiate representation. As we defined the term, something counts as a representation when it has the function of being reliably caused by something external to it. Representations, on this view, are something like tracking devices. Tracking, it turns out, is no trivial matter. The things we track tend to vary from one context to the next. Therefore, a tracking device is successful only if it can adapt to these changes. Tracking is necessary for representation, and sensitivity to context is necessary for tracking. It follows that context sensitivity is necessary for representation, not incompatible with it.

EMBODIMENT

Traditional theories in cognitive science adopt representations that are inherently disembodied. A disembodied representation is one harbored by a system that lacks a sensory-motor system capable of interacting with the environment. As a result, disembodied representations do not rely on sensory-motor systems to do their work—instead they rely solely on their intrinsic representational resources. By focusing on disembodied representations as paradigms for exploring cognition, cognitive scientists have failed to develop theories that are well suited for actual performance in the real world. Because disembodied representations typically take the form of arbitrary language-like codes, nothing inherent in their structure supports grounding them in bodies and environments.

We agree with dynamicists that disembodiment is a fatal flaw for any representational theory and that embodied forms of these theories cry out for development. We further agree with dynamicists that brain–body–environment interactions are central to many, if not all, intelligent functions. However, we do not believe that embracing embodiment implies that dynamic systems must be the correct account of cognition and that representational approaches must be incorrect. On the one hand, representational theories can be made inherently embodied, as we demonstrate later in presenting perceptual symbol systems. On the other hand, dynamic theories are not necessarily embodied. Indeed, disembodied dynamic theories are widespread in cognitive science.

Embodiment in Connectionist Systems

As we have already observed, some dynamicists do not repudiate representations. Instead, they try to appropriate them in a dynamic framework. Connectionists typically fall into this camp. Many of them have even tried to meet Fodor and Pylyshyn's (1988) demand for representations that can be productively combined (van Gelder, 1990). For example, Pollack (1990) and Smolensky (1990) have developed distributed connectionist systems that superimpose vectors for finite elements to productively create complex representations. Because such systems can apparently construct infinite numbers of complex representations from finite elements, they are productive. Furthermore, they seem to improve on classic computational accounts of productivity by implementing elemental concepts in dynamic and context-sensitive manners. As the same concept is superimposed in different complex representations, it adopts context-sensitive forms that have dynamic properties. These approaches openly admit that cognitive systems are representational, but they identify representations with activation patterns in artificial neural networks. So construed, representations can be analyzed by using the tools of dynamic systems theory.

Although these connectionist representations implement productivity in a context-sensitive manner, they are inherently disembodied. The vectors that represent elemental concepts are essentially arbitrary strings that bear no systematic relation to their perceptual and motor referents. To see this, consider a standard feed-forward connectionist net that implements completely distributed representations in a single layer of hidden units. As Barsalou (in press) discussed, the layer of input units is often interpreted as a perceptual system that detects the presence or absence of sensory features. In contrast, the layer of hidden units is often viewed as providing conceptual interpretations of inputs. Most critically, the mapping from input units to hidden units is arbitrary and reflects the small random weights assigned to the connections between them before learning. Depending on different random assignments, the mappings from input to hidden units vary arbitrarily. As a result, conceptual interpretations in a feed-forward net are disembodied, just like those in standard computational representations. In both cases, conceptual interpretations have nothing in common with perceptual states, with the two having an arbitrary relation.

As this example illustrates, dynamic approaches to cognition are not necessarily embodied, although they can be. In this regard, dynamic approaches do not differ from their more traditional counterparts. As we show shortly in our discussion of perceptual symbol systems, it is feasible to develop accounts of productive representation that are context sensitive *and* embodied.

Embodiment in Situated Robotics

Researchers in situated robotics have been more fervently committed to embodiment than have other researchers. Exciting developments in this area have exposed serious problems with classical theories. To be effective, the cognitive system that guides a robot's behavior must interface well with perception and action. A system that simply computes abstract functions does not suffice. Instead, a system must be able to use sensory information effectively and to implement action effectively. Whereas cognitive systems inspired by traditional computational theories have fared poorly in these regards, cognitive systems inspired by dynamic systems have fared much better (Brooks, 1991).

Perhaps the most important lesson learned from these recent developments is that using the resources of a cognitive system solely to accomplish perception and action is misguided. It may well be the case that no formulation of a cognitive system in isolation can ever accomplish these tasks on its own. Instead, a much more successful tack is to use resources inherent in the body and the environment. Rather than relying solely on cognitive mechanisms, intelligence evolves from dynamic interactions between brains, bodies, and environments (Clark, 1997a, 1997b; Thelen & Smith, 1994). All three components are essential, and no one alone suffices.

To accomplish these dynamic interactions, dynamicists often call for the elimination of representations altogether. However, the representations that inspire these calls are the classical representations of traditional cognitive science, namely, representations that are static, context independent, and disembodied. We have already argued that advanced intelligent systems simply cannot do without representations. In our discussion of context sensitivity, we began to argue that important alternatives exist to classical representations, alternatives that are compatible with dynamicists' aims. A similar moral applies with regard to embodiment. Rather than dispensing with representation, we should look for forms of representation that are more intimately connected to sensory-motor systems, which mediate our interaction with the world. We will pursue this possibility more fully when we discuss perceptual symbol systems below. The interim moral is that, on the one hand, connectionist nets do not guarantee embodiment, and, on the other hand, situated robots try to do without representation. We want to steer a course between these options by adopting embodied representations.

TEMPORALITY

Of all the features emphasized by proponents of dynamic systems theory, temporaility is claimed to be the most important (van Gelder, 1997). Traditional approaches in cognitive science are atemporal. We define a cognitive model as atemporal if it does not explicitly define states or state transitions with reference to real time. In contrast, the differential equations used by dynamic systems theorists typically include variables for time. In so doing, they can capture temporal distances between events, rates of change, and duration of processes. Traditional approaches often use classical computational models, which can explain such phenomena only in terms of computational steps. Such steps are defined without reference to real time, so that they fail to capture truly temporal properties of a system. Dynamic models are to be applauded for their ability to capture such properties. Their introduction into cognitive science promises to shed light on phenomena that cannot be readily modeled with more traditional tools. Some dynamicists seem to endorse a stronger claim, however. They think we can do away with models that fail to explicitly capture temporal properties. These include most standard representational models. We are uncomfortable with this move.

Abstracting Over Real Time

In sharp contrast to the dynamicists, we think the ability of traditional representational theories to abstract over real time is one of their strongest virtues. By ignoring the spatiotemporal details of neural activity, traditional representational models identify the stable, functional properties of cognition. On these approaches, cognitive activity can be characterized in terms of the representations that are used and the processes that govern them without reference to how long such processes take. In this manner, representational models establish the high-level, functional properties of cognition, while ignoring the low-level details of implementation. Indeed, one can infer from the absence of real time in many traditional models that timing belongs to a noncognitive level of analysis.

To bolster the intuition behind this approach, consider the following thought experiment. Imagine Slow World, a place exactly like Earth except that everything takes twice as long. Each of us has a doppelgänger on Slow World. Of course, the snail-paced speeds of people on Slow World do not impair them, because everything else in their world occurs at the same slow pace. In fact, if we were to view a sped-up video recording of Slow World, it would be indistinguishable from our own. Our doppelgängers would do just what we do in comparable situations (compare Block, 1978).

This comparability suggests that our Slow twins should be classified with us from the perspective of psychology. They have minds just like ours and can be subsumed under a common psychological theory. However, if psychological laws defined cognition only in terms of real time, this commonality would not be apparent. Laws that solely specify the real-time properties of Earth cognition would be inapplicable to Slow World cognition. Laws focusing only on real time would be inadequate to capture the important similarities between these two forms of cognition. Instead, establishing an account of cognition that captures these similarities requires an atemporal analysis—exactly the type of analysis that traditional representational models provide. Traditional models that abstract over real time can exactly establish those *relative* temporal properties we share with our Slow World doppelgängers.

Dynamicists might respond in several ways to this core assumption of the traditional representational approach. First, they might argue that the obvious similarities between individuals in the two worlds can be captured by saying that they share similar dynamic systems, with one running at twice the rate of the other. To specify how these dynamic systems are similar, however, one must abstract from their temporal properties, which concedes that there is a useful level of explanation that ignores real time. Taking a more offensive tack, dynamicists might respond by arguing that ignoring the dimension of real time leads to the omission of critical mechanisms. In particular, a traditional account of people in the two worlds we are considering fails to capture the very temporal properties that distinguish them. This response raises a valid point, but it only points to the compatibility of dynamic and atemporal approaches. We need the former to explain the differences between Slow World and the real world, and we need the latter to explain the similarities.

Implementing Real-Time Constraints

One can underscore the importance of real-time constraints by considering the 100-step rule of Feldman and Ballard (1982). Using physiological facts about the rate of neural firing, Feldman and Ballard arrived at a conservative estimate of how long a neural-processing step takes in real time. By establishing the time to perform some task and dividing it by this estimate, one arrives at the maximum number of steps an account of that task can include. Of course, this analysis makes assumptions about the seriality of steps, which might not hold in a cascade model (McClelland, 1979). Nevertheless, the point remains that real time must play a critical role in evaluating cognitive theory. Theories of cognition that do not address real time not only fail to account for important phenomena, but may also be implausible once the implications of real-time processing are considered.

At first blush, this seems to be an argument against cognitive models that do not define operations in terms of real time. On closer analysis, however, the lesson of Feldman and Ballard is that we should construct models that can be implemented in a way that respects real time regardless of whether they are temporal models. Numerous atemporal models appear capable of implementing real-time performance. For decades, traditional computational models in mainstream psychology have been developed to explain all sorts of real-time processing, including activation, matching, search, retrieval, transformation, and response generation. So many models exist that trying to cite them all here takes more space than allowed for this chapter (for a few examples, see Just & Carpenter, 1992; Logan, 1988, 1996; McClelland, 1979; Nosofsky & Palmeri, 1997; Ratcliff, 1978; Townsend, 1990; Van Zandt & Ratcliff, 1995). Our present point is not to argue that classical computational models are adequate, but to argue that adequate models can abstract away from real time. It is simply not true that atemporal models cannot naturally and elegantly implement real-time processing. Put differently, a model that is atemporal in the sense of lacking explicit real-time variables can be temporal in the sense of satisfying real-time constraints. So conceived, atemporality and temporality are perfectly compatible.

Dynamicists might be inclined to reformulate their temporal challenge in another way. Cognitive models that abstract from real time are typically discrete. Discrete models, it is argued, cannot explain the continuous properties of cognitive systems (van Gelder, 1995). Most obviously, discrete models fail to capture the fact that natural cognitive systems reside in some state at every moment of continuous time. As a result, such models cannot provide adequate accounts of cognition.

There are a number of replies to this objection. The first we owe to Eric Dietrich, who pointed out (personal communication) that the continuity assumption may concede too much to proponents of dynamic systems theory. Nature may ultimately bottom out into discrete interactions between discrete particles. In that case, apparent continuity rests on a discontinuous foundation. We are also grateful to Dietrich for suggesting a delightful thought experiment. Imagine Strobe World, a place that appears much like our own world, but where people and objects flicker very quietly in and out of existence. Strobe World is certainly discrete, but we would not deny that its inhabitants think.

The continuity challenge can be met even if we concede that real cognitive systems are continuous. Fifty years ago, Alan Turing (1950) anticipated this challenge and responded as follows:

> Strictly speaking, there are no [discrete state] machines. Everything moves continuously. But there are many kinds of machines that can be profitably *thought of* as being discrete machines. For instance in considering the switches for a lighting system it is a convenient fiction that each switch must be definitely on or off. (p. 441)

As this passage illustrates, when discrete models are *implemented,* they are discrete only in a fictitious sense, because their physical states change continuously over time. It is simply false that such models are inherently noncontinuous. An implemented model can occupy different states at every point in time. In some cases, these differences do not make a difference. State individuation can abstract from minute changes if these play no explanatory role from the point of view of psychology. On this approach, the fiction of discreteness is profitable, because continuity plays no theoretically interesting role.

Adopting another strategy, proponents of computational approaches can admit that continuous properties in real time are theoretically significant, while insisting that atemporal theories constitute useful approximations. In general, it is always possible to provide a continuous interpretation of a discrete model. For example, when computers compute integrals, they sometimes approximate them in discrete steps at the programming level. Nevertheless, the users of these computations often interpret them as reflecting continuous mathematical processes. Similarly, when cognitive scientists implement activation functions in computational models (e.g., Anderson, 1976), they implement them discretely at the programming level but interpret them continuously at the theoretical level. As these examples illustrate, it is readily possible to conceive of computational mechanisms as having continuous properties, even though they are discrete at the programming level. At the physical level, they once again become continuous, but typically in ways that differ considerably from their conceptions at the theoretical level (e.g., the continuous transition of states in a silicon chip does not correspond to the continuous transitions in a theoretical activation function). Again, our aim is not to defend computational approaches, but to demonstrate that atemporal models can be defended against the charge that they fail to capture continuity. In some cases continuity does not make a theoretical difference, and, in others, discrete models provide useful approximations of theoretically interesting continuities.

It must be reiterated that we are not arguing against models that make reference to real time. Cognitive systems have real-temporal properties that are interesting and important. We only suggest that admitting the value of such models poses no threat to models that abstract from real time. If they are capable of satisfying real-time constraints, such models may be adequate and useful for certain explanatory purposes. There is no reason to think that cognitive models must explicitly use real-time variables.

There is a further argument that concedes this point, but continues to insist that the accommodation of temporal properties poses a threat to traditional representational models. According to this line, traditional representations are too cumbersome or inefficient to satisfy real-time constraints. Natural cognitive systems negotiate complex environments with remarkable speed and do so under considerable temporal pressure. For example, if one encounters a hungry tiger in the wild, one had better be able to quickly come up with an escape strategy. Too much cogitation is deadly. Those interested in situated cognition have taken such examples to heart. A growing number of researches believe that we do not achieve behavioral success by internally representing the situations that confront us. Brooks (1991), for example, has repudiated representation on the grounds that it takes too much time. Meeting real-world demands does not give us the luxury of representation.

We think this line of argument is partially correct. Time constraints have implications for what kinds of representations situated agents can deploy. They do not, however, show that no representations are deployed. This point is taken up again in the next section after we introduce perceptual symbol systems. Perceptual symbol systems show that a representational account of cognition can explain context sensitivity and embodiment while satisfying real-time constraints.

PERCEPTUAL SYMBOL SYSTEMS

We began our discussion by arguing that theories of cognitive systems should not dispense with internal representations. We then argued that representational accounts could be context sensitive, embodied, and accordant with real-time constraints. Along the way, we have suggested that such accounts can be formulated without identifying representations with constructs introduced by the tools of dynamic systems theory. At the same time, we have been happy to admit that such tools may be fruitfully employed to describe cognitive systems at nonrepresentational levels of analysis. To illustrate that representational systems can be context sensitive, embodied, and temporal, we present *perceptual symbol systems* as a theory of cognition. Because we develop detailed accounts of this approach elsewhere (Barsalou, in press; Prinz, 1997), we do not do so here. Instead, we simply provide a brief summary emphasizing the features most relevant to this discussion.

The first premise of perceptual symbol systems is that cognition and perception share common representational mechanisms, namely, the sensory-motor areas of the brain. This is analogous to the argument that imagery and perception share common neural bases on at least the visual, motor, and auditory modalities (e.g., Crammond, 1997; Farah, 1995; Jeannerod, 1995; Kosslyn, 1994; Zatorre, Halpern, Perry, Meyer, & Evans, 1996). Neuroscientists who study category-specific losses have reached a similar conclusion about category knowledge, namely, that it is grounded in the sensory-motor areas of the brain that process its referents (e.g., Damasio, 1989; Damasio & Damasio, 1994; Gainotti, Silveri, Danieli, & Giustolisi, 1995; Rösler, Heil, & Hennighausen, 1995). Intuitively, the idea is that conceptually representing a category involves running perceptual systems as if one were actually perceiving a member of that category. Thus, conceptually representing chairs involves running visual, haptic, motor, and somatosensory systems as if they were actually interacting with a physical chair. Of course, such simulations are not necessarily identical to the neural states that underlie perception, but they are cut from the same cloth, running the same systems in a similar manner.

Most important, this view of representation departs considerably from the disembodied views that permeate traditional computational theory. Rather than using amodal symbols that lie outside perceptual areas to represent knowledge, perceptual symbol systems use representations in these areas. As a result, they are inherently embodied.

Perceptual symbol systems are also inherently representational, because their primary purpose is to create representations that both detect physical objects, when used in recognition, and "stand in" for physical situations, when run offline (Clark, 1997a, 1997b; Donald, 1991, 1993). Furthermore, perceptual symbols, like classical representations, are capable of combing productively and representing propositions. Productivity results from the ability to compose perceptual simulations hierarchically and recursively. Propositions result from binding simulations to perceived individuals. We do not develop the details of these accounts here (see

Barsalou, in press; Prinz, 1997, 1998). Instead, we simply note these abilities to make the point that perceptual symbol systems are inherently representational, exhibiting classic representational functions.

Perceptual symbol systems must be context sensitive to achieve simulation competence (Barsalou, in press). Essentially, a simulation competence is a set of mechanisms that allow an individual to competently simulate the members of a category in their absence. The simulation competence for a particular category has two levels of structure. First, it contains a tremendous amount of perceptual category knowledge that has accumulated over a lifetime of experience with category members. For example, the simulation competence for *car* contains all the multimodal, sensory-motor knowledge that someone has acquired about cars. Second, a simulation competence produces specific simulations (mental models) of the category on specific occasions. For example, the simulation competence for *car* can produce an infinite number of specific simulations, thereby representing many different kinds of cars in many different situations. Besides being able to represent actual cars previously seen, a simulation competence can represent novel cars and schematic cars, using productive abilities to construct simulations.

Context sensitivity arises naturally from simulation competence. Rather than constructing a single simulation of *car* to represent the category in all contexts, a specific simulation is tailored to each one. Thus, representing cars on the French Riviera may produce simulations of expensive European cars, whereas representing cars on a hippie commune may produce simulations of painted Volkswagen buses. Similarly, representing cars being driven may produce simulations from the driver's perspective, whereas representing cars being repaired may produce simulations of a car's mechanical systems. In this manner, a simulation competence implements a wide variety of context-sensitive phenomena, including framing effects, context-dependent features, and representational flexibility between and in subjects (Barsalou, in press).

Perceptual symbols systems also satisfy real-time constraints. This is best demonstrated by considering the Brooks (1991) objection that we glossed at the end of the last section. According to Brooks, cognitive systems ought to dispense with internal representations, because forming such representations precludes real-time success. When acting under time pressure, it is more efficient to bypass a representational stage and let sensory inputs issue directly into motor outputs. We think there is a very important kernel of truth to this argument, but it also makes an instructive mistake. A system that goes directly from sensory inputs to motor outputs is not, thereby, nonrepresentational. The sensory inputs and the motor commands can themselves qualify as representations. Such a system lacks only *amodal* representations. The perceptual symbol systems view repudiates amodal representations. It agrees with Brooks that postulating an amodal level of representation is dangerously inefficient. The remaining sensory-motor states, however, qualify as representational. Like traditional representations, they carry information about the world, and they can be run offline in planning. Unlike traditional representations, they bypass a costly transduction into an amodal code and are thereby more capable of satisfying real-time constraints.

Perceptual symbol systems exhibit another interesting temporal property. Because conceptual knowledge involves running perceptual and motor systems in roughly the same manner as in perception and action, conceptual knowledge re-enacts, at least somewhat, the temporal properties of these original experiences. Representing a category conceptually plays out temporally in at least somewhat the same manner as experiencing category members. As a consequence of this view, certain representations are conceived as temporally extended sequences. A representation of an event is not a static description represented as a single processing step, but a temporally sequenced series of representations corresponding to the activities that make up that event. The theory itself does not need to introduce a real-time variable to describe such sequences, but the sequences would have a real-temporal profile when implemented. In this way, perceptual symbol systems make predictions about real-time performance.

As this brief summary illustrates, perceptual symbol systems attempt to reinvent the strengths of traditional approaches and dynamic systems approaches in a single theory. On the one hand, perceptual systems are inherently representational; on the other, they are inherently embodied, context sensitive, and temporal.

COMPLEMENTARY APPROACHES TO MODELING COGNITION

Perceptual symbol systems are not formulated by using the tools of dynamic systems theory. A perceptual symbol is characterized as a facsimile of a perceptual state that can stand in for external entities in offline simulations. We have not identified perceptual symbols with state-space trajectories, attractors, or other dynamic constructs. At the same time, we are open to the possibility that such identifications may be appropriate at another level of analysis. More generally, we believe that dynamic tools can be used in conjunction with perceptual symbols theory in providing a complete account of cognition. The two approaches have distinct, but complementary explanatory virtues. Dynamic systems provide a natural account of continuously varying processes and real-time profiles. Perceptual symbols theory provides a natural account of more stable properties such as internal representations. The best account of cognition would come from merging these resources.

Such a merger can be achieved in different ways. Obviously, the two approaches can be used to model different levels. Dynamic systems are well suited for capturing low-level neural events (as in Freeman's work). They may also be well suited for capturing a very high level of analysis, such as the movement of an organism toward a food source (Busemeyer & Townsend, 1995). In contrast, perceptual symbols can capture the stabilities that supervene on low-level states or the representational mechanisms that guide an organism's dynamic actions.

Perceptual symbols and dynamic systems can also be compatibly used to explain different phenomena. Dynamic systems can be fruitfully used in explaining environmentally responsive motor movement (Thelen & Smith, 1994). They can also be used to model developmental changes over time. For example, the notion of a bifurcation has been used to explain qualitative transitions without appeal to preordained developmental stages (Elman et al., 1996). Perceptual symbols, in contrast, may be better suited for explaining more cognitive phenomena such as categorization, analogical reasoning, and planning.

Perhaps the most intriguing possibility for demonstrating compatibility is to integrate perceptual symbols and dynamic systems into a unified model of a single phenomenon. For example, both may together explain context sensitivity. Earlier, we argued that a representational account can accommodate context sensitivity by forming novel combinations from a relatively stable feature set. The category of dogs, for example, may be represented by different perceptual symbols on different occasions. The feature selection process, however, may be usefully explained by postulating dynamic mechanisms that capture the rising and falling of features and the constraint satisfaction that converges on a final set. An integration of this kind may provide a more complete account of this process than an account stated in terms of either approach alone.

In our opinion, it is overly simplistic and limiting to argue that either dynamic or representational systems provide the best approach to modeling cognition. Although each approach captures important insights, neither captures everything that must be explained. The two can be integrated by modeling different levels of analysis, modeling different phenomena, or forming hybrid models of the same phenomena. The perceptual symbols approach invites integration with dynamic approaches, because it departs from the context-insensitive, disembodied, and static representations favored by more traditional theories. These shared objectives suggest that the two approaches may be complementary.

Clearly, the construction of unified theories constitutes an open frontier. We offer perceptual symbol systems as one preliminary attempt in this spirit, yet we are under no illusions about its current status. This approach obviously requires significant development before it satisfies the challenges we raise here. Much must be accomplished before its implementations of embodiment, context sensitivity, and temporal processing are adequately formulated and defended. Nevertheless, we believe that this sort of theory exemplifies the theory of the future. We should be attempting to develop theories that reinvent previous approaches, salvaging their strengths, and discarding their weaknesses. Rather than disavowing internal representations, we should reconstrue them in a way that accords with the demands of embodied, real-world cognition. In this endeavor, perceptual symbol systems and dynamic systems approaches can form a natural alliance.

ACKNOWLEDGMENTS

We thank our editors, Eric Dietrich and Arthur Markman, for their helpful comments on an earlier draft.

REFERENCES

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory and Cognition, 10,* 82–93.

Barsalou, L. W. (1987). The instability of graded structure in concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101–140). New York: Cambridge University Press.

Barsalou, L. W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76–121). New York: Cambridge University Press.

Barsalou, L. W. (1990). On the indistinguisability of exemplar memory and abstraction in category representation. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 3, pp. 61–88). Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of memories* (pp. 29–101). London: Lawrence Erlbaum Associates.

Barsalou, L. W. (in press). Perceptual symbol systems. *Behavioral and Brain Sciences.*

Bechtel, W. (1996, July). Yet another revolution? Defusing the dynamic systems theorists' attack on mental representations. Presidential address to the Society for Philosophy and Psychology, San Francisco.

Block, N. (1978). Troubles with functionalism. In C. W. Savage (Ed.), *Minnesota studies in the philosophy of science, vol. ix: Perception and cognition: Issues in the foundations of psychology* (pp. 261–362). Minneapolis: University of Minnesota Press.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence, 47,* 139–160.

Busemeyer, J., & Townsend, J. (1995). Dynamic representation of decision making. In R. F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 101–120). Cambridge, MA: MIT Press.

Chomsky, N. (1957). *Syntactic structures.* The Hague, Netherlands: Mouton.

Clark, A. (1997a). *Being there: Putting brain, body, and world together again.* Cambridge, MA: MIT Press.

Clark, A. (1997b). The dynamic challenge. *Cognitive Science, 21,* 461–481.

Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese, 101,* 401–431.

Crammond, D. J. (1997). Motor imagery: Never in your wildest dreams. *Trends in Neuroscience, 20,* 54–57.

Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the neurosciences, 2,* 263–275.

Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition, 33,* 25–62.

Damasio, A. R., & Damasio, H. (1994). Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 61–74). Cambridge, MA: MIT Press.

Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition.* Cambridge, MA: Harvard University Press.

Donald, M. (1993). Precis of "Origins of the modern mind: Three stages in the evolution of culture and cognition." *Behavioral and Brain Sciences, 16,* 739–791.

Dretske, F. I. (1981). *Knowledge and the flow of information.* Cambridge, MA: MIT Press.

Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes.* Cambridge, MA: MIT Press.

Dretske, F. I. (1995). *Naturalizing the mind.* Cambridge, MA: MIT Press.

Edelman, G. M. (1989). *The remembered present: A biological theory of consciousness.* New York: Basic Books.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: MIT Press.

Farah, M. (1995). The neural bases of mental imagery. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 963–975). Cambridge, MA: MIT Press.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science, 6,* 205–254.

Fodor, J. A. (1990). *A theory of content and other essays.* Cambridge, MA: MIT Press.

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

Freeman, W. J. (1991, February). The physiology of perception. *Scientific American, 264,* 78–85.

Freeman, W. J., & Skarda, C. A. (1990). Representations: Who needs them? In J. L. McGaugh, N. M. Weinberger, & G. Lynch (Eds.), *Third conference, brain organization and memory: Cells, systems and circuits* (pp. 375–380). New York: Guilford Press.

Gainotti, G., Silveri, M. C., Daniele, A., & Giustolisi, L. (1995). Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory, 3,* 247–264.

Giunti, M. (1995). Dynamic models of cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 549–571). Cambridge, MA: MIT Press.

Jeannerod, M. (1995). Mental imagery in the motor context. *Neuropsychologia, 33,* 1419–1432.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99,* 122–149.

Kosslyn, S. M. (1994). *Image and brain.* Cambridge, MA: MIT Press.

Logan, G. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

Logan, G. (1996). The CODE theory of visual attention: An integration of space-based and object-based attention. *Psychological Review, 103,* 603–649.

McClelland, J. L. (1979). On the time-relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86,* 287–330.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100,* 254–278.

Mundale, J., & Bechtel, W. (1996). Integrating neuroscience, psychology, and evolutionary biology through a teleological conception of function. *Minds and Machines, 6,* 481–505.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266–300.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

Pollack, J. (1990). Recursive distributed representations. *Artificial Intelligence, 46,* 77–105.

Prinz, J. J. (1997). *Perceptual cognition: An essay on the semantics of thought.* Unpublished doctoral dissertation, University of Chicago.

Prinz, J. J. (1998). *Regaining composure: A defense of prototype compositionality.* Manuscript submitted for publication.

Rösler, F., Heil, M., & Hennighausen, E. (1995). Distinct cortical activation patterns during long-term memory retrieval of verbal, spatial, and color information. *Journal of Cognitive Neuroscience, 7,* 51–65.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence, 46,* 159–216.

Stich, S., & Warfield, T. A. (Eds.). (1994). *Mental representation: A reader.* Oxford: Blackwell.

Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: MIT Press.

Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science, 1,* 46–54.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59,* 433–466.

Van Essen, D. C., Felleman, D. J., DeYoe, E. A., Olavarria, J., & Knierim, J. (1990). Modular and hierarchical organization of extrastriate visual cortex in the macaque monkey. *Cold Spring Harbor Symposia on Quantitative Biology, 55,* 679–696.

van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science, 14,* 355–384.

van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy, 92,* 345–381.

van Gelder, T. (1997). Dynamics and cognition. In J. Haugeland (Ed.), *Mind design II* (pp. 421–450). Cambridge, MA: MIT Press.

Van Zandt, T., & Ratcliff, R. (1995). Statistical mimicking of reaction time: Single-process models, parameter variability, and mixtures. *Psychological Bulletin and Review, 2,* 20–54.

Zatorre, R. J., Halpern, A. R., Perry, D. W., Meyer, E., & Evans, A. C. (1996). Hearing in the mind's ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience, 8,* 29–46.

Zeki, S. (1992, September). The visual image in mind and brain. *Scientific American, 267,* 68–76.