Cognitive Psychology **40,** 138–171 (2000)
doi:10.1006/cogp.1999.0727, available online at http://www.idealibrary.com on IDEAL®

In Defense of Representation

Arthur B. Markman

University of Texas

and

Eric Dietrich

Program in Philosophy, Computers, and Cognitive Science,
Binghamton University

The computational paradigm, which has dominated psychology and artificial intelligence since the cognitive revolution, has been a source of intense debate. Recently, several cognitive scientists have argued against this paradigm, not by objecting to computation, but rather by objecting to the notion of representation. Our analysis of these objections reveals that it is not the notion of representation per se that is causing the problem, but rather specific properties of representations as they are used in various psychological theories. Our analysis suggests that all theorists accept the idea that cognitive processing involves internal information-carrying states that mediate cognitive processing. These mediating states are a superordinate category of representations. We discuss five properties that can be added to mediating states and examine their importance in various cognitive models. Finally, three methodological lessons are drawn from our analysis and discussion. © 2000 Academic Press

INTRODUCTION

Representations have been a critical explanatory tool in cognitive science. Virtually all theories about cognition are based on hypotheses that posit men-

This work was supported by NSF CAREER award no. SBR-95-10924 to the first author. An earlier version of this paper appeared in the electronic journal *Psycoloquy*. The authors thank Mark Bickhard, Hugh Clapin, Art Glenberg, Stevan Harnad, Pat Hayes, Eliot Hearst, Dedre Gentner, Scott Kelso, David Krantz, Gary Marcus, Len Matin, Clay Morrison, Janet Metcalfe, Tomislav Pavlicic, Robert Remez, Herb Terrace, Esther Thelen, and Takashi Yamauchi for helpful comments during the evolution of this project. Thanks also to John Elway for 16 years of raw strength and courage and Jayson Williams for great hustle on the court.

Address correspondence and reprint requests to Arthur B. Markman, Department of Psychology, University of Texas, Mezes Hall 330, Austin, TX 78712. E-mail: markman@psy.utexas.edu.

0010-0285/00 $35.00

tal representations as carriers of information about the environment of the organism or agent. Recently, however, researchers have argued that the value of representations in cognitive science has been exaggerated (e.g.,Brooks, 1991; Thelen & Smith, 1994; van Gelder & Port, 1995). Many of these researchers argue that representations should be eliminated from cognitive models, and in their place psychology should focus on the relationship between the cognitive system and the environment or on the sub-representational dynamics of cognitive systems. These views take representations to be at best, emergent entities from more basic dynamics and, at worst, constructs that have stunted the growth of cognitive science.

In fact, a review of the anti-representationalist claims in the literature reveals that researchers do not seem to be objecting to representation so much as objecting to specific properties of representations that seem obligatory in the context of modern cognitive science. Indeed, despite the disagreement among researchers, virtually everyone in the field of psychology and artificial intelligence (AI) seems to agree on the existence of some sort of information-carrying state internal to a cognitive system as well as on the need for these internal states in cognitive theories.

The problem with this dispute over representation is that it is divisive and makes it difficult to resolve key theoretical issues about what kinds of properties mental representations have. In this article, we provide a view of representation designed to move the debate over representation out of its current morass into deeper issues about the properties of representation necessary to explain various cognitive capacities. We do this by first identifying a core notion of representation that we believe is acceptable both to traditional cognitive scientists as well as to those who have called on us to abandon representation. We call this core a *mediating state.* It is possible to add a variety of additional properties to this core in order to generate representations that have been proposed by various cognitive scientists, and we consider five such properties in this article, namely (1) being enduring, (2) being discrete, (3) having compositional structure, (4) being abstract, and (5) being rule governed. Finally, at the end of the article we present three suggestions for how representation should be used in cognitive science.[1]

THE CONCEPT OF A MEDIATING STATE

In this section, we define the central construct of a mediating state. This construct captures the core idea of a representation and provides a structure

---

[1] We also need to say what we don’t do. We do not provide a theory of representational content or semantics. That is far beyond the scope of this project. Representational content is an important topic (see Bickhard & Terveen, 1995), but it is not the only topic worthy of discussion. Figuring out what is right about representation, and what all the new debate about the utility of representations means, is also important. This is what we do here.

for considering additional properties of representations. The basic idea is that mediating states are internal states of a *system* that carry *information* which is used by the system in the furtherance of its *goals*. (Not all system states are information states; some are goal states.)

Mediating states are the superordinate category of representations. That is, mediating states are to representations what birds are to eagles. Presenting this superordinate category is important because people have rejected the notion of representation in cognitive science. If there were ornithologists going around saying there were no eagles, then a good strategy for setting up a debate between those ornithologists who did and those who did not believe in eagles would be to point out that there are birds and, furthermore, that they *agreed* there are birds. Then, the debate could focus on the properties of existing birds to determine whether any of these birds are recognizable as eagles. We adopt a similar strategy here.

Our discussion of mediating states is rather detailed because a significant reason for the debate over representation is that there is no generally accepted definition for the concept. Thus, in many arguments about representation, the opposing sides seem to focus on different aspects of representation (as an example, see the paper by Vera and Simon, 1993, as well as the responses to it). By providing a common framework for talking about representation, we aim to avoid this problem.

To begin, we consider the general idea of a cognitive system. Any entity that exercises some control on its environment via feedback loops must have internal states for comparing the actual states of its environment with ‘‘desired’’ states. We call the desired states *goal states* or just *goals*.[2] We call the internal states which denote the actual states of the environment (within a certain degree of accuracy) *information states*. We say that the internal information states contain *information*. We call such information states *mediating states*. Goals have long been recognized as important components of the cognitive system (e.g., Lewin, 1935) and some computational modelers suggest internal states are necessary for systems to pursue goals (Agre, 1995).

We define a ‘‘system’’ as any entity that uses information in an attempt to satisfy its goals. Systems, on this definition, must have feedback loops (at least negative ones) because they need to determine whether their goals are satisfied. The goals of the system might only be implicit, i.e., not explicitly represented anywhere in the system. A thermostat controlling a heater in an enclosed building is a good example of our definition of a system. The

---

[2] This might seem overly broad, but it is not. Most of the time, goals further the *life* (or continued existence) of the entity, but this need not be the case. Consider a robot that is exploring Mars, but cannot never return or a human on some crucial, sacrificial mission (like a terrorist with bombs attached to his body). Both have goals, but neither has goals for staying alive.

goals of the thermostat system are not in any sense known to it. On our view, systems are capable of making errors but not necessarily of correcting them. For example, one can hold a match under a thermostat and get it to behave as if the room temperature were quite high, though the temperature in the room might be below freezing.

All systems, on our definition, have and use information. Thus, we need some definition of information that is sufficient to support our definition of a mediating state. At the outset, we must be clear that we are not providing a definition of how representations come to have semantic content. This problem has vexed philosophers for a long time, and we do not solve it here. Instead, we provide a definition of the informational content of a mediating state. It is likely that such information figures into a theory of semantic content, so we use a definition of information that is likely to be compatible with future theorizing about semantic content.

To this end, we use a version of a definition of information presented by Dretske (1981) modified in several ways for use in psychological explanation. Dretske's concept of information was derived from Shannon's (1948) quantitative definition of information. For Shannon, information is measured as the average amount of data arriving at a receiver, generated by a source, and transmitted across a channel from the source to the receiver. The problem with this idea is that it provides no way of considering the information content of a specific signal. Instead, it considers only the amount of information, averaged over all possible signals. Dretske altered Shannon's definition to permit the informational content of a specific transmitted signal from source to receiver to be considered. He defined information in terms of conditional probabilities: an event, e, at a receiver carries the information that something, s, has property P, $P(s)$, if and only if the conditional probability of $P(s)$ given e is 1, $\mathrm{Prob}[P(s)|e] = 1$. [The same conditional probability can also be used to talk about the signal, r, causing e as the carrier of the information that $P(s)$.] We consider the receiver to be the system in question and the information channel (the channel over which a signal is transmitted) to be the lawful connection between the energy type(s) the system is sensitive to, the thing giving off (or reflecting) the energy, and the system's sensory apparatus. Because the system has goals, it can and does affect its environment. These effects are further sources of information transmitted to the receiver (e.g., telling the system that it has achieved a goal).

Now we must deal with a complication. We cannot use Dretske's definition as it stands because it implies (at least on the surface) that the cognitive system cannot be in error. Dretske's definition of informational content relies on the fact that the conditional probability that something, s, has property P is 1 given that some cognitive event, e.g., a perceptual judgment that $P(s)$, has occurred. In symbols:

$$\mathrm{Prob}[P(s)|J(P(s))] = 1.$$

This requirement is a problem because it is possible that s is not P. For example, suppose you wake up in the middle of the night and see what looks like your dog in your room, and indeed form the perceptual judgment that your dog is lying there. In the morning you discover that what you took to be your dog is a pile of rumpled clothes. If the conditional probability is 1 that there is a dog lying there given that you've judged that the object is your dog, then your dog had to have been in your room. Hence there is no room for error on Dretske's definition of information. Thus, this definition is too strong.

There are several technical solutions to this problem. One possibility is to give up the idea that information is important to the content of a mediating state; that is, we could abandon Dretske's definition of information. This solution is unappealing because the idea that information is important to content is really just a formal way of saying that there must be a causal link between an organism and its environment. Given the lack of success of any purely causal account of content, the best way to satisfy the intuition that causation is important to content is to assume that information (as we defined it) is important to content.

A second possibility is to relax the requirement that $\text{Prob}[P(s)|e] = 1$. This option is available because (unlike Dretske) we are not interested in making information the sole basis of the semantic content of mediating states. We use something like this below for higher order mediating states, but this solution won't work for low-level mediating states. For low-level mediating states, we turn to a third option, which is motivated by an example.

Consider the psychology of mistaking a pile of clothes for a dog. Presumably, in the prior history of a cognitive system that can recognize dogs, it has learned a set of perceptual features that reliably indicate the presence of a dog. Enough of these features are activated by looking at the pile of clothes that the system reaches the judgment that a dog is present. Some of the features are quite specific and are tied to the perceptual environment (e.g., the presence of particular edges, as detected in the pattern of light on the retina). Other features are more abstract and may arise as a function of many different possible patterns of sensory stimulation (e.g., ears or a snout). The more abstract features are crucial for object recognition because they allow a person to recognize dogs they have never seen before (see Biederman, 1987, for a similar discussion). On this view, these abstract features are reliable, but not perfect, indicators of the presence of a dog. This is because the connection between low level perceptual features and the presence of these more abstract features in the environment is imperfect and merely reliable (i.e., the probability of the abstract features given the low level perceptual features is less than 1). Further, most theorists in categorization believe that there are no necessary and sufficient features that determine membership in a category (e.g., Smith & Medin, 1981). Thus, the probability that an object is a member of a particular category given a set of features (even when the presence of those features in the environment is known with certainty) is less than 1. Thus, the nature of the conceptual system introduces the possibility of error in category judgments in order to allow it to recognize novel instances in the world.

This point is important, so we say it in another way. Features used for recognition are mediating states, but these features are *not* the perceptual judgment that there is a dog present; rather, they are the low-level building blocks of such perceptual judgments. Our proposal is that a conditional probability equal to 1 should be associated with the activation of these features (these mediating states), rather than with the high-level perceptual judgment itself, and that the higher level mediating states, such as perceptual judgments and the like, inherit the information only with some level of reliability, i.e., higher–level mediating states inherit the relevant information with a probability less than 1. At this time, we do not have a fully worked out theory of how this information inheritance works, but we have enough apparatus here to push forward with our discussion of mediating states.

Now we can say that the conditional probability begins as 1, thereby tying the content of all mediating states to information without having to *identify* that content with the information. More formally, let $D^* = \text{DF1, DF2, . . ., DF}n$ be a set of (mental) low-level features active in some system and sufficient for recognizing a dog. These may only be a subset of all the features that can be used to recognize a dog, but let us assume that they work reliably. Let $E^* = \text{EC1, EC2, . . ., EC}n$ be environmental conditions, whatever those are, that cause the dog features to become activated. Then the following is true:

$$\text{Prob}[E^*|D^*] = 1. \qquad (1)$$

It *is* true that there are conditions in the subject's environment that are sufficient for activating the subject's low-level dog features because, in fact, those initial features got activated. They must have been activated some way. And the only way to do it is via some environmental conditions. Normally, activating those features leads to the correct conclusion that a dog is present via the activation of some higher mediating states (e.g., the concept for ''ears,'' ''snout,'' and finally ''dog''). But those features can be falsely activated by a rumpled pile of clothes in dim light. In the latter case, the high-level perceptual judgment that there is a dog present is wrong because the inheritance relationship is not perfectly reliable, but the low-level features nevertheless got activated for the right reasons, namely, there were the right low level conditions in the environment.

This now gives us a multitier approach to representational error: the low-level mediating states are not in error, but higher level states introduce error in the process of inheriting the relevant information and making categorical judgments. Thus, additional levels of mediating states may provide a road toward a solution to the problem of misrepresentation while still tying information to semantic content without making information the whole of semantic content. (Briefly, what we believe is needed is the idea of *functional role*. Higher level mediating states inherit the initially accurate low-level information only with some probability less than 1. What keeps the higher level mediating states synchronized with the environment and with each other is their functional role with respect to each other.)

In sum, on our view, mediating state information is always *about* something, typically something removed in space and time from the system. For example, in the thermostat–room–heater system, the room temperature co-varies with the curvature of the bimetallic strip. Thus, there is information in the system which is used to affect the environment external to it. We call states of information of this type *mediating states*.

Now that we have described the sort of informational states that we are interested in, we can define a mediating state in terms of the following four necessary and jointly sufficient conditions.

(1) There is some entity with internal states which include goal states; we assume that these states undergo changes.
(2) There is an environment external to the system which also changes states.
(3) There is a set of informational relations (which at the low levels is accurate and grades off somewhat at the higher levels due to imperfect inheritance) between states in the environment and the states internal to the system. The information must flow both ways, from the environment into the system and from the system out to the environment. (In the simplest case, this is a feedback loop, but more complicated loops such as plan–act–detect loops are also possible. Note also that in the typical case, these informational relations are realized as causal relations, but what is important is the information carried by these causal relations, not the causal relations themselves.)

If we stopped here, there would be a problem (just as there is with Dretske's view of information) because there is no way to distinguish the information content of the mental states of a cognitive system from those of other states such as transducer states (e.g., an eardrum) or index states (e.g., a sunburn). Thus, we add the constraint that the information must be used by the system in achieving its goals.

(4) The system must have internal processes that act on and are influenced by the internal states and their changes, among other things. These processes allow the system to satisfy system-dependent goals (though these goals need not be known explicitly by the system).

We close this section by making several observations. First, this definition of a mediating state is quite general. It is intended to capture something that all cognitive scientists can agree to; namely, that there is *internal* information used by systems (organisms, for example) that mediates between environmental information coming in and behavior going out (this is the minimal condition that distinguishes cognitive science from behaviorism). Interestingly, most computational models in psychology and AI do not use actual mediating states because their internal states do not actually bear any correspondence to relevant entities outside the system. This result is sufficient to show that the definition of mediating state we have derived is far from vacuous: the data structures of the vast majority of AI systems do not satisfy it. Perhaps one prominent exception is robots like those developed by Brooks (1991), which seem to have rudimentary mediating states that link them to their environment (we describe these kinds of robots in more detail below). Nonetheless, the absence of true mediating states has not prevented AI systems from being useful both as tools and as explanatory models.

From the point of view of mediating states, there is more agreement than disagreement among representationalists and antirepresentationalists. We argue that disagreements over whether there are representations are in fact better and more usefully understood as different researchers focusing on different aspects of cognition and using different kinds of mediating states to explain what they are observing. These different kinds of mediating states are obtained by adding one or more of the five properties of mediating states described below. Also, from this perspective, one can see that there is an important but underappreciated diversity among research strategies and representational methodologies.

Mediating states capture the core of what is important about representations from an explanatory point of view. The five key properties of mediating states mentioned above (and discussed in the next section), either singly or jointly, are essential to explaining many kinds of cognitive processes. Some researchers might want to call mediating states with some or all of the five properties added on representations. If so, that is fine. But no theoretical point turns on this. The debate in cognitive science should not be over whether representations left in their intuitive sense are necessary, but rather over which particular properties of mediating states are necessary to explain particular cognitive processes. As we elaborate below, this view suggests that cognitive science should strive for a diversity of research methodologies that brings to light explanatorily useful properties of mediating states rather than being antirepresentationalist or seeking the one true representational formalism that serves as the basis of all cognitive processes.

Finally, it is important to emphasize again that the definition of mediating states is not intended to function as a theory of representational content. Such

TABLE 1
Five Properties That Can Be Added to Mediating States, but Have Been Rejected by Some Antirepresentationalists

1. Cognitive systems have no enduring representational states.
2. Cognitive systems have no discrete symbols.
3. Cognitive systems have no compositional structures. They do not permit role–argument bindings.
4. Cognitive systems do not have abstract representations.
5. Cognitive systems do not operate via rule-governed processes.

a theory is very much needed in cognitive science, but there is as yet no consensus on the details of such a theory.[3]

FIVE PROPERTIES THAT CAN BE ADDED TO MEDIATING STATES

This section examines five properties that mediating states can have. Each of these properties has been the subject of an antirepresentationalist argument. These properties are (1) being enduring, (2) being discrete, (3) having compositional structure, (4) being abstract, and (5) being rule governed (see Table 1).

Are There Enduring Mediating States?

Many attacks on traditional representational assumptions have focused on the fluidity of cognitive processing (Thelen & Smith, 1994; van Gelder & Port, 1995). These attacks contrast this fluidity with the rigidity of classical cognitive models. It is argued that traditional models are not sufficiently flexible to account for the fine details of cognitive processing. In place of traditional representational systems, critics posit cognitive models that in-

[3] It cannot be overstressed that conditions 1–4 are *not* conditions which specify how mediating states (or representations) get or have semantic content. It is not part of the definition of mediating states that what they mean is the same thing as what they carry information about. In particular, we are not claiming that information plays the role of *fixing* semantic content. Many have misread our four conditions this way. We think there is a lot more to representational content than the information the representation carries.

We are specifying what mediating states are by specifying how they function in cognition. As we said, mediating states are abstractions of representations. As an analogy, Gregor Mendel specified the notion a *gene* without specifying the actual ‘‘content’’ of genes, i.e., without specifying the sequences of DNA that comprise genes. It is true that he put broad constraints on what the ‘‘content’’ of genes could be, but he did not specify genetic content in any detail at all. We are doing the same thing for mediating states: we are providing existence conditions for such states without specifying how they have the content that allows them to do what they do. As we said, the issue of semantic content in mediating states (and in representations) is important to psychology and cognitive science, but it is not the only issue. Agreeing that there are such things as mediating states and hence representations seems prior to developing a theory of their content.

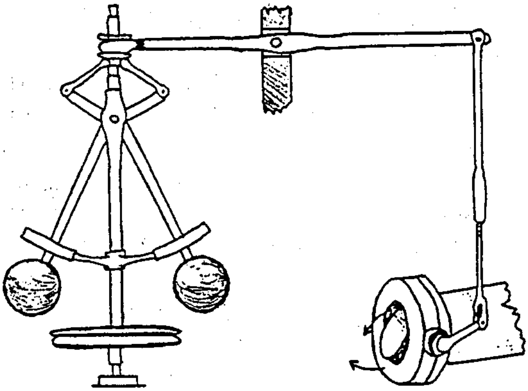

FIG. 1. Watt's steam engine governor (source; Thelen & Smith, 1994).

volve dynamic moment-by-moment changes in the internal states of the system.

One prominent example of such a system is Watt's steam engine governor (Thelen & Smith, 1994; van Gelder & Port, 1995). The steam engine governor is a remarkable device designed to keep steam engines from exploding. The governor is attached to a steam pipe. As the engine runs faster, the governor spins faster causing balls mounted on arms on the either side of the governor to rise. The rising balls cause a valve to close because of mechanical connections between the arms and the valve. The constricted valve decreases the amount of steam flowing, and hence the pressure, which causes the engine to slow, which in turn causes the governor to spin more slowly. The slower spin causes the balls on the arms to drop, opening the valve. In this way, a relatively constant pressure inside the governor can be maintained (see Fig. 1).

This example has been used to demonstrate that interesting behavior can be carried out without representation. The power of the example rests on the fact that the mediating states of the governor are not enduring. Instead, the speed of the governor at a given moment is an indirect function of the pressure in the engine at that moment (because the speed of the engine is related to the pressure). When the pressure changes, the speed of the governor changes. However, there is no record of past states of the system. It is suggested that, just as the steam engine governor does not need enduring mediating states, cognitive systems do not need them either.

Transient mediating states are not limited to simple mechanical objects like governors or thermostats. Kelso (1995) points out that a horse's gait is controlled by an elaborate dynamic interaction that involves the speed of the limbs at that moment, but is not influenced by previous gaits (except for the gait at which the horse is currently moving). The patterns of activation on units in distributed connectionist models are also transient states. When a new pattern of activity arises on a set of units (perhaps on the output units as the response to a new input), the old pattern is gone. Similarly, the current state of a dynamic system is transient and changes to some new state as the system evolves.

Although dynamic systems clearly involve transient changes in the activation of their mediating states, they also require states that endure over longer periods of time. Connectionist models use the weights on the connections between units as a trace of past activity. Of course, these mediating states are highly distributed. No particular weight (or unit) can be identified as a symbol in presentation, although the behavior of a connectionist network can often be examined by looking at properties of the connection matrices, such as their principal component structure. These connection weights are enduring mediating states; without such states, connectionist models could not learn. In general, dynamic systems must have some enduring energy landscape that determines how a new state can be derived from an existing one. This landscape determines key aspects of the behavior of the system, such as the location of attractor states.

There is an appealing insight in the view that cognitive systems do not have enduring mediating states; namely, that not all behaviors that involve mediating states require these states to be enduring. For example, the classic studies of the gill withdrawal reflex in the sea slug *Aplysia* have demonstrated that this reflex can be habituated with repeated stimulation. Kandel and his colleagues (Klein, Shapiro, & Kandel, 1980) demonstrated that with repeated stimulation of the gill, the presynaptic motor neuron in the circuit involved in the habituation releases less neurotransmitter into the synaptic cleft than it did initially. This decrease in transmitter release appears to be mediated by a blockage of calcium channels in the presynaptic neuron. When stimulation of the gill is terminated, the calcium channels become unblocked and the gill withdrawal reflex returns to its original strength. In this system, the amount of transmitter released into the cleft is a mediating state that controls the desired strength of the gill withdrawal reflex, which is translated into an actual strength of the reflex by the postsynaptic motor neuron. This state is not enduring, however. With repeated gill stimulation, more calcium channels become blocked or, conversely, as the habituation stimulus is extinguished, more channels unblock. In either case, the reflex only has a current level of activity. It does not store past states.

This discussion demonstrates that some mediating states are transient. Nonetheless, not all mediating states can be transient. Rather, systems that learn and make use of prior behavior have some enduring states that allow the system to react to new situations on the basis of past experience. A steam engine governor or a sea slug gill may operate on the basis of transient states, but enduring mediating states are required in models of more complex cognitive behaviors (Agre, 1995).

Are There Discrete Mediating States?

Many cognitive models assume that the elements in a representation are discrete and composable. (Discrete states are referred to as ‘‘entities’’ to emphasize their discreteness. In the literature, such entities are frequently called ‘‘symbols,’’ but this term is not used here because it is question begging.) Discrete entities are elements in many proposals about cognitive representations, ranging from feature-list representations (Chomsky & Halle, 1968; Tversky, 1977) to semantic networks (Anderson, 1983b; Collins & Loftus, 1975) to structured representations and schemas (Norman & Rumelhart, 1975; Schank & Abelson, 1977). Despite the variety of proposals that cognition involves discrete entities, there have been many arguments that such entities fail to capture key aspects of cognitive processing.

It might seem that being a discrete composable entity follows directly from being an enduring representational entity, but not all enduring mediating states make finite, localizable, and precise contributions to larger states. For example, attractor states in dynamic systems and in iterative connectionist models are enduring, but they are not discrete entities (or symbols). Attractor states are not clearly separable from each other by distinct boundaries; hence their semantic interpretations are not precise. In these systems, transitions from one state to the next occur as the result of processes like energy minimization that operate over continuous states.

The idea that there are no discrete entities in cognitive systems reflects the important insight that new cognitive states are never (or almost never) exact duplicates of past ones. New states may bear some likeness to past states, but they are not identical. In a distributed connectionist model this insight is reflected in the idea that new states are activation vectors that are similar to (i.e., have a high dot-product with) activation vectors that have appeared in past states. In a dynamic systems model, a cognitive system whose behavior is characterized as a point in a state space may occupy points in neighboring regions of state space without occupying the same point more than once.

Smolensky (1988; 1991) made this point explicitly in his defense of connectionist models. He proposed that subsymbolic representations could be used to model the effects of context on cognitive processing. For example, a discrete symbol for ‘‘cup’’ captures very little information about cups. Rather, the information about cups that is relevant to a cognitive system changes with context. For thinking about a cup full of coffee, insulating properties of cups are important, and examples of cups that have handles may be highly accessible. For thinking about a ceremonial cup, its materials and design may be more important than its insulating properties. Smolensky argued that the high degree of context sensitivity in cognitive processing militates against discrete entities as the basis of cognitive states.

Clark (1993) raises a related question about where discrete entities come from. He suggests that connectionist models might actually be better suited to cognition than classic symbolic processes because their sensitivity to statistical regularities in the input may help them develop robust internal states that still have most of the desirable properties discrete entities are supposed to provide.

One suggestion for how context sensitive representations might be acquired was presented by Landauer and Dumais (1997). They describe a model of the lexicon (Latent Semantic Indexing, LSA) that stores higher order co-occurrence relations among words in a sentence using a high-dimensional space (e.g., one with 300 dimensions). One interesting property of this model is that its performance on vocabulary tests improves both for words seen in the passages presented to it on a given training epoch and for words that were not seen during that training epoch. This improvement on words not seen is due to the general differentiation of the semantic space that occurs as new passages are presented. Despite its excellent performance on vocabulary tests [when trained on encyclopedia articles, LSA performs the TOEFL (Test of English as a Foreign Language) synonyms test at about the level of a foreign speaker of English], it contains no discrete entities corresponding to elements of word meaning.

A second line of research that poses a problem for systems with discrete entities focuses on the metacognitive feelings engendered by cognitive processing (Metcalfe & Shimamura, 1994; Reder,1996). For example, we often have a ‘‘feeling of knowing.’’ When we are asked a hard question, we might not be able to access the answer to it, but we may be quite accurate at saying whether we would recognize the answer if we saw it. This feeling seems to be based on the overall familiarity of the retrieval cue (Metcalfe, 1993; Reder & Ritter, 1992) as well as on partial information retrieved from memory (Koriat, 1994). Neither of these processes seems to involve access to discrete properties of the items being processed.

Despite the evidence for continuous mediating states, there are some good reasons why complex cognitive systems must have mediating states with discrete parts. There is evidence that when people make comparisons among concepts, their commonalities and differences become available (Gentner & Markman, 1997; Markman & Gentner, 1993; Tversky, 1977). For example, when comparing a car and a motorcycle, people find it easy to list commonalities (e.g., both have wheels; both have engines) as well as differences (e.g., cars have four wheels, motorcycles have two wheels; cars have bigger engines than motorcycles). If a model has discrete entities, then it is possible to access those entities later. In contrast, if a model does not have discrete entities, then the individual parts of a representation cannot be accessed and used by cognitive processes. For example, in a distributed connectionist model, the active mediating state at a given time consists of a pattern of activity across a set of units. Typically, processing involves comparing vectors using a holistic strategy like the dot product, which calculates the amount of one vector that projects on another. A scalar quantity like the dot product loses all information about what aspects of one vector are similar to another, yielding only a degree of similarity. Only when there are discrete entities can there be access to the content of the commonalities and the differences.

A similar problem arises for Landauer and Dumais's high dimensional semantic space model described above. As discussed , this model performs well on the synonyms test from the TOEFL by finding words near to it in semantic space (in this case by having a high dot-product). Its success on this test is offered as evidence of its adequacy as a model of human lexical processing that does not require discrete entities. However, this system would have difficulty with an antonyms test. Antonyms are also words that are highly related to each other, but differ along a salient dimension (e.g., ''up'' and ''down'' differ in direction, and ''up'' and ''down'' are more similar to each other than either is to ''giraffe''). Selecting the word most similar to the target would likely find synonyms, but finding the word most dissimilar to the target would find unrelated words. Determining the antonym of a word requires analyzing the parts of the relevant lexical mediating state, and these parts are simply not available in a purely high-dimensional semantic space.

Another reason discrete entities seem important for cognitive processing comes from studies demonstrating that people can (depending on the circumstance) have a preference for or against exact matches along a dimension. In the study of similarity, Tversky and Gati (1982) found that people tend to give high weight to identity matches (see also Smith, 1989). In contrast to the predictions of mental space models of mental representation, Tversky and Gati found that pairs of stimuli that could each be described by values on two dimensions were considered more similar when one of the dimensions for each stimulus was an exact match than when both dimensions had similar but not identical values. Interestingly, the opposite result has been found in studies of choice (Kaplan & Medin, 1997; Simonson, 1989). When faced with a choice, people often select an option that is a compromise between extreme values. For example, an ideal diet meal might be one that tastes good and has very few calories. People on a diet given a choice among (1) a meal that tastes good and has many calories, (2) a meal that tastes fair and has a moderate number of calories, and (3) a meal that tastes bad, and has very few calories are likely to select the middle option (2) because it forms a compromise between the extremes. In this case, the exact match to an ideal is foregone in favor of an option that partially satisfies multiple active goals. In these examples, the objects have a part identity rather than an overall identity. It is not clear how a system without discrete entities would find pairs with some identical aspects to be so compelling.

To summarize, not all cognitive processes require mediating states with discrete elements. Dynamic systems and connectionist models that use spatial mediating states are often good models of cognitive behavior. These processes may often be sensitive to context. Nonetheless, the influence of context can also be modeled with discrete mediating states that have a small grainsize. Other processes, such as finding antonyms or making comparisons, seem to require at least some mediating states that are discrete.

Are There Mediating States with Compositional Structure?

An important observation about cognitive processing is that concepts combine. This ability to form more complex concepts from primitive units is particularly evident in language, where actions are described by the juxtaposition of morphological units that represent objects (typically nouns) with other units that represent relations between those objects (typically verbs). Because we combine concepts freely and easily in this manner, it is often assumed that symbolic representations have a compositional (or ‘‘role argument’’) structure that facilitates combination (e.g., Fodor & McLaughlin, 1990; Fodor & Pylyshyn, 1988).

A central problem with cognitive processes that require a role–argument structure is that they require processes that are sensitive to the bindings between predicates and their arguments. Structure-sensitive processes are often much more complex than processes that can operate on non–compositional structures (or states). For example, when a mediating state is spatial, processing involves measuring distance in space (like the dot product in connectionist models). When structures are independent symbols (or sub–symbolic features), then sets of features can be compared using elementary set operations (as in Tversky’s 1977 contrast model). However, when structures have bindings, a compositional procedure that is sensitive to those bindings must be created. Often, the processes proposed by cognitive scientists have been quite complex.

Consider the act of comparing two structures, as in the example of comparing the atom to the solar system. One popular model of comparison, Gentner’s (1983, 1989) structure-mapping theory, suggests that comparisons seek structurally consistent matches, meaning that the match must obey both parallel connectivity, and one-to-one mapping. In parallel connectivity for each matching predicate, the arguments to those predicates must also match (e.g. the electrons corresponds to the planets because both are revolving around something). One-to-one mapping requires that each element in one structure match at most one element in the other (e.g., mapping the electrons to the planets means they cannot also correspond to the sun). Thus, the comparison process takes into account the bindings between predicates and their arguments. A number of computational procedures for determining analogical matches have been developed (Falkenhainer, Forbus, & Gentner, 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997; Keane, Ledgeway, & Duff, 1994).

While it may be appropriate to assume that some cognitive processes have

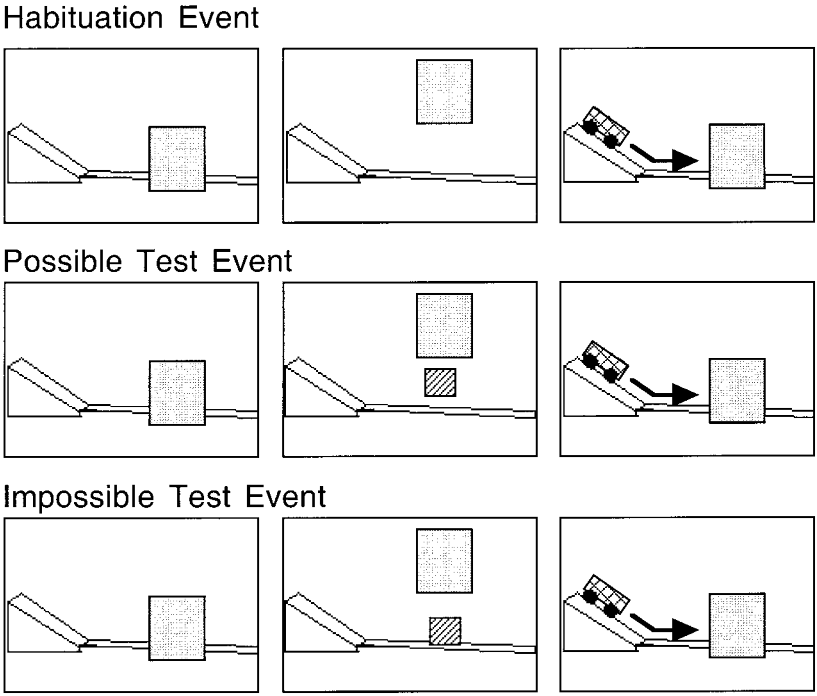

FIG 2. The design of the object permanence tasks designed by Baillargeon.

such complexity, it has been suggested that structure-sensitive processes are inappropriate as models of cognitive development. Indeed, a central problem that Thelen and Smith (1994) raise with the representational view of mind is that it posits representations and processes that seem far more complex than make sense on a developmental account. As one way to address this point, they discuss explanations for Baillargeon's (1987) classic studies demonstrating that infants have object permanence (see Fig. 2).

As shown in the top row of Fig. 2, infants are habituated to an event in which a screen is lowered and then a car on a track rolls down a ramp, goes behind the screen, and reemerges on the other side. This task is presented repeatedly until the infant's looking time to this event subsides. Then, both possible and impossible test events are presented. In the possible event (shown in the middle row), a block sits behind the track, the screen lowers, and the car again rolls down the ramp to behind the screen and emerges on the other side. In the impossible event (shown in the bottom row), a block sits *on* the track, the screen lowers, and the car rolls down the ramp, goes behind the screen, and emerges on the other side. Infants show greater looking time to the impossible event than to the possible one. This result is interpreted as a recognition that the block continues to exist behind the screen and should have stopped the progress of the car.

An explanation of this event involving a compositional symbol system would assume that infants store specific relationships such as that the block was *on* the track or the block was *behind* the track as well that the car was *on* the track. It is critical to this explanation that the child can make a distinction between the consequences of the block being *on* the track and the block being *behind* the track. The process underlying this behavior might be specific to cars and tracks (or perceptual objects of particular types) or it might be general to moving objects and obstructions.

Thelen and Smith suggest that this explanation grants too much knowledge to an infant. In particular, they reason that if infants could form a symbolic representation of the scene, then it is not clear why they should require a sequence of habituation trials in order to form their representation. Moreover, if infants have such elaborate knowledge of objects, it is not clear why they should act as if hidden objects did not exist in traditional Piagetian object permanence tasks. Thus, Thelen and Smith suggest that the symbolic account of this task provides a gross description of infants' behavior, but fails to explain the details.

In place of a symbolic model, Thelen and Smith propose a dynamic systems account. They suggest that the infant reacts to regularities detected by the visual system. The infant visual system is assumed to have systems that specify *what* objects exist in the world and *where* those objects are located. These outputs form a state space. The impact of habituation is to form an expected trajectory through the state space. Then, during the test events, Thelen and Smith assume, the child dishabituates to the impossible event because its trajectory starts out similar to that of the habituation event but then diverges from it at some point. In contrast, the trajectory of the possible event does not diverge from that of the habituation event, so no dishabituation is observed.

This dynamic systems account is intriguing, but we suggest that it cannot explain the infants' behavior without positing a complex compositional structure—a mediating state with a role–argument structural description. In the highly impoverished form of the ''what'' and ''where'' systems in the example, it is not clear what information is supposed to be captured in the visual array. However, even if a complex array of values were sufficient to model the output of these systems, there is no account of why the trajectory divergence caused by having a block *on* the track is more surprising than the trajectory divergence caused by having the block *behind* the track. That is, Thelen and Smith provide no account of how an undifferentiated notion of trajectories in a state space distinguishes between trajectory differences that matter and those that do not. We suggest that infants' behavior in this case must reflect a recognition of the spatial relationships between objects and that augmenting the dynamical systems view to account for these data ultimately requires the addition of a capacity for storing discrete compositional spatial relations. That is, mediating states with a role–argument structure are needed.

A brief examination of research in visual object recognition suggests that visual mediating states may be profitably characterized as having components that encode spatial relations between parts. Kosslyn (1994) marshals behavioral, computational, and neuropsychological evidence in favor of the hypothesis that there are two different modes that the visual system uses to describe relationships between elements in images. The right hemisphere system describes the visual world in terms of metric aspects, and the left hemisphere system uses qualitative relations between elements to describe the world (although see Ivry & Robertson, 1998, for an alternative explanation of these findings). Other behavioral and computational evidence that visual object recognition requires attention to relations between parts in images comes from Biederman (1987; Hummel & Biederman, 1992) and Palmer (1977). For example, Biederman (1987) suggests that mediating states denoting objects consist of primitive shapes connected by spatial relations [see also Marr (1982)]. As evidence, he demonstrates that the ability to recognize objects in line drawings is disrupted more by eliminating information at the junctions of line segments (which carries information about relations between parts) than by eliminating an equivalent amount of line information between the joints. This work further suggests that the visual array required by Thelen and Smith's explanation of the object permanence studies is likely to involve some relational elements. This interpretation is reinforced by the observation that spatial prepositions refer to spatial relations that are abstracted away from many specific details of objects (e.g., Herskovits, 1986; Landau & Jackendoff, 1993; Regier, 1996).

Finally, as discussed at the beginning of this section, compositional structure seems necessary for models of linguistic competence. Many linguists and philosophers have pointed out that people effortlessly distinguish between sentences like ‘‘The Giants beat the Jets’’ and ‘‘The Jets beat the Giants.’’ Even connectionist models of phenomena like this make use of structured internal states (e.g., Chalmers, 1990; Elman, 1990; Pollack, 1990). We are also able to keep track of others' beliefs when they are explicitly stated. Thus, I may believe that the Giants beat the Jets last week, but that you believe the opposite. A ‘‘propositional attitude’’ like belief requires not only that I be able to encode the elements in the original proposition itself (that the Giants beat the Jets), but that you believe the opposite proposition (so I must be able to represent the metaproposition). This sort of processing admittedly requires effort (see Keysar, Ginzel, & Bazerman 1995) and does not develop immediately (Perner, 1991; Wellman, 1990), but it eventually becomes a significant part of human linguistic competence. It seems unlikely that these abilities could be modeled without mediating states that have a role–argument structure.

In sum, one insight underlying the proposal that representations do not have compositional structures is that such representations require significant effort to construct and also significant effort to process. This complexity seems to go beyond what is required to carryout many cognitive tasks. It would be an overgeneralization, however, to conclude that compositional structures are not needed at all. Tasks as basic as those that demonstrate object permanence in infants and processes like-object recognition clearly involve at least rudimentary relations between objects in a domain. A model that has no capacity for role–argument binding cannot explain the complexity of such higher level cognitive and linguistic processing.

Are There Abstract Mediating States?

A common intuition is that abstract thought is central to cognitive processing. At one level, it is trivially true that mediating states are abstract. The world itself does not enter into our brains and affect behavior. Even sense data are the result of neural transformations of physical stimuli that reach on our sense organs. Hence, the question being raised is more accurately cast as a search for the level of abstraction that characterizes mediating states. The classic assumption is that the information we store is extremely abstract and hence that it applies across domains. Indeed, when a logical statement like $P \rightarrow Q$ is written, it is assumed that any thinkable thought can play the role of P or Q. This assumption has been called into question.

One source of the attack on highly abstract stored information comes from demonstrations that people's performance on logical reasoning tasks is often quite poor. For example, in the classic Wason selection task (Wason & Johnson-Laird, 1972), people are told to assume that they are looking at a set of four cards that all have a number on one side and a letter on the other and that they must select the smallest set of cards they would have to turn over in order to test the rule ''If there is a vowel on one side of the card, then there is an odd number on the other side.'' The four cards show an A, 4, 7, and J, respectively (see Fig. 3). In this task, people appear sensitive to the

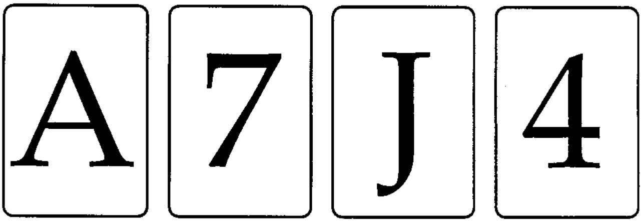

FIG. 3. The Wason selection task: if there is a vowel on one side of the card, there is an odd number on the other side (source, Markman, 1999).

logical schema called modus ponens ($\mathrm{P} \rightarrow \mathrm{Q}, \mathrm{P} :: \mathrm{Q}$), as virtually all people state that the card with the A on it must be turned over. In contrast, people generally seem insensitive to modus tollens ($\mathrm{P} \rightarrow \mathrm{Q}, \sim\mathrm{Q} :: \sim\mathrm{P}$), as few people suggest that the card with the even number must be turned over. Further support for this finding comes from studies of syllogistic reasoning in which people exhibit systematic errors in their ability to identify the valid conclusions that follow from a pair of premises (Johnson-Laird, 1983).

These errors have been explained by appealing to abstract logical rules that differ in their ease of acquisition (Rips, 1994). However, much work has focused on more content-based structures that might be used to solve logical problems. For example, Johnson-Laird and his colleagues (Johnson-Laird, 1983; Johnson-Laird, Byrne, & Tabossi, 1989; Johnson-Laird & Byrne, 1991) have suggested that people solve logical reasoning problems by constructing mental models that contain familiar objects and use inference rules derived from familiar situations. Consistent with this claim, it has been demonstrated that people's performance on logical-reasoning tasks like the Wason selection task is much better when the situation is specific and familiar than when it is abstract. Studies have demonstrated that people perform well on the Wason task when the scenario involves social rules like permission, obligation, or catching cheaters (Cheng & Holyoak, 1989; Cosmides, 1989). Although debate continues over the exact nature of people's reasoning processes, there is general agreement that content has a strong influence on how people reason.

The content-bound nature of reasoning has led some researchers to assume that the bulk of human reasoning is inseparable from the world. The robots developed by Brooks (1991) embody this assumption. Brooks's robots do not form extensive structures to describe their environments; they only use information that is immediately available and store only transient information as they navigate the world. Modules in the robot communicate with each other only by allowing one module to inhibit the activity of another without passing any information between them.

A related approach is taken in psychology in the study of situated action (Clancey, 1997). For example, Hutchins (1995) performed a far-reaching study of navigators aboard naval ships. He argues that the complex task of plotting a course for a ship involves deep cognitive work by (at least some) of the participants, but it also requires extensive use of tools and of shared information processing. No individual has an abstract structure of the entire navigation task. Instead, the task itself is structured by the tools used to complete it (such as maps and protractors).

Cognitive linguists have also taken the view that mental structures are not entirely abstract. Langacker (1986) suggests that syntactic structures reflect states of the world. The encoding of prepositions like ''above'' and ''below'' are assumed to be tied to structures that encode spatial information rather than simply reflect abstract structures. The linguistic representations are symbolic, but they are assumed to be symbols that are closely tied to perceptual aspects of the world. This contrasts with the amodal verbal symbols often used in linguistic models. Thus, cognitive linguistics assumes a much closer connection between syntax and semantics than does classic linguistics.

Mainstream research in cognitive science has also shifted away from the use of abstract logical forms toward more content-based approaches. In the study of categorization, significant progress has been made by assuming that people store specific episodes rather than abstractions of category structure (Barsalou, 1999; Brooks, 1978; Medin & Schaffer,1978; Nosofsky, 1986). Research on problem solving has demonstrated that people solve new problems by analogy with previously encountered problems rather than on the basis of abstracted solution procedures (Bassok, Chase, & Martin, 1998; Novick, 1990; Reed & Bolstad, 1991; Ross, 1984). For example, Bassok et al. (1998) found that arithmetic word problems written by college undergraduates were affected by the content of the word problems. If arithmetic knowledge were truly abstract, then these content effects would not be expected. In AI, the field of case-based reasoning has taken as a fundamental assumption that it is easier to store retrieve and tweak existing cases than to form abstract rules, derive procedures for recognizing when they should be used and then adapt them to be applied in a new situation (Kolodner, 1993; Schank, Kass, & Riesbeck, 1994).

These examples demonstrate that there are unlikely to be general-purpose context-free schemas and processes that are ready to be deployed in whatever domain they are needed. The fact that cognitive processing generally shows strong effects of content means only that most mediating states contain some information about the context in which they were formed. It does not mean that there is no highly abstract information stored in some mediating state somewhere. It is likely that the information within an individual may differ in its degree of abstractness. Some types of inference schemas (like modus ponens) seem so obvious and independent of the domain that we may very well store them as abstract rules (see Rips, 1994, for a similar discussion). Other types of inferences seem to rely heavily on the domain. The main question to be answered by cognitive science is how many kinds of mediating states are abstract and how many are concrete and what level of abstraction is used by different cognitive processes. Currently, the balance seems to favor concreteness for many cognitive processes.

Are Mediating States Rule-Governed?

Many classic models of cognitive processing posit rules. For example, Piagetian stage theory is based on the idea that the end state of development is a competence with formal operations. Chomskian linguistics assumes that grammar involves rules that transform a deep structure into a surface structure. Classic AI assumes that problems can be solved by applying operators in which the presence of a set of antecedent conditions cause a consequent that changes values of mediating states and controls effectors that interact with the world.

In classic AI reasoning systems, the rules are generally inference schemas or productions (Anderson, 1983b, 1993; Newell, 1990; Pollack, 1994). The rules may also be statistical procedures that are supposed to capture crucial elements of expert reasoning behavior. In AI research on problem solving, a problem is cast as a discrepancy between a beginning and an end state, and problem solvers are assumed to have an array of rules (or operators) that can be applied that reduce this discrepancy (Newell & Simon, 1963). On this view, problem solving is a search through a problem space generated by the application of rules to the current state.

A rule-governed approach is also evident in developmental psychology. As Smith and Sera (1992) point out, many developmental theories begin with the adult behavior as the expected end-state and then develop a theory that leads inexorably from an inchoate beginning state to an orderly adult competence. Adult behavior is often described in terms of a system of rules and children are then monitored until they show sensitivity to the proper set of adult rules. For example, in Piagetian studies of the balance beam, children are given long blocks of various shapes and encouraged to try to balance them on a fulcrum. They are monitored for the development of the correct rule that the downward force of a weight is a function of the weight and the distance of the weight from the fulcrum. In this task, children's behavior is often described as the development of intermediate (and incorrect) rules like ''the fulcrum must always be in the center.'' On this view, developmental milestones consist of the acquisition of particular rules.

In many ways, these models of development resemble linguistic models. A central tenet of modern linguistics is that syntactic structure is guided by a highly abstract and universal set of rules determining which sentences are grammatical in a given language. Linguistics is concerned primarily with linguistic competence—an accurate description of the grammar of a given language. Psychologists who have adopted this framework (and have studied linguistic performance) have assumed that there is some mental representation of these syntactic structures. On this view, the sentences of a language are constructed through the application of grammatical rules. Many psycholinguistic models posit processes in which rules are applied to linguistic input that allow the structure of the sentence to be determined from its surface form.

Antirepresentationalist arguments have often centered on rules. A central argument by Thelen and Smith (1994) is that cognitive development does not involve the acquisition of rules. They use the development of locomotor ability as an example. As they point out, very young infants exhibit a stepping motion when their feet are stimulated if their weight is supported externally. This ability later seems to disappear, only to reemerge still later in development. Many theories of locomotion used this description of the behavior of the average child as the basis of theories of motor development. These theories often posit maturational changes that permit the observed behaviors to occur.

Thelen and Smith (1994, Thelen, 1995) argue that the rule-based view does not properly characterize children's development. Children supported in water exhibit the same stepping behavior as younger infants supported out of water, leading to the conclusion that increases in the weight of the legs may be causing the observed cessation of stepping behavior. Support for this comes from studies in which leg weights are attached to very young infants, which causes the stepping movements to stop. The fine details of locomotor behavior suggest that children's development is guided not by the acquisition of a small set of rules, but rather by the interaction of multiple physical and neural constraints. Behavior is guided in part by the maturation of brain and tissue. It is also guided by a child's interaction with the outside world. A variety of factors must come together to shape development. Finally, the rule-based view of motor development focuses on the progression of the average child and ignores individual differences. In contrast, the dynamic view of motor development considers individual variation to be important data (see also Kelso, 1995).

These examples provide compelling evidence that rules are not needed in explanations of many cognitive processes. Many systems can be described by rules, but that is not the same as using rules to carry out a process. For example, the steam-engine governor has a mediating state (the speed with which the governor spins), and a mechanism that makes use of that state (a combination of arms and levers that close the valve as the height of the arms increases). The system is not checking the state of memory in order to determine the appropriateness of a rule. Thus, the system is not actually using a rule to carry out its behavior.

Although we agree that many cognitive processes do not need rules, it does not follow that rules are not a part of any cognitive system. Indeed, some cognitive processes seem like good candidates for being rule-based systems. For example, there have been no convincing accounts to date that the statistical structure of a child's linguistic input is sufficient to lead them to acquire a grammatical system consistent with their language. Furthermore, there have been some impressive demonstrations of rule use. For example, Kim, Pinker, Prince, and Prasada (1991) demonstrated that verbs derived from nouns are given a regular past-tense form, even when the noun is the same as a verb that takes an irregular past-tense form. So in describing a baseball game an announcer will say ''Johnson flied out to center field his first time up'' rather than ''Johnson 'flew' out . . .'' because the verb ''flied'' is derived from the noun fly [ball] rather than the verb to fly. Furthermore, Marcus, Brinkmann, Clahsen, Wiese, and Pinker (1995) have demonstrated that the way a verb is given its past tense form or a noun its plural form need not be a function of the frequency of that morphological ending or the similarity of the verb or noun to known verbs and nouns in the language. These findings support the view that grammar is mediated by rule-governed processes.

This discussion of rules requires one technical point and one methodological point related to it. If the computational hypothesis about the nature of cognition is correct [and it is a hypothesis, not a loose metaphor (Dietrich, 1990, 1994, p.15)], then it *must* be possible in principle to model cognition using rules because it is a theorem in computability theory that a rule-based machine can do everything a Turing machine can do. Put another way, if cognition involves the execution of algorithms, then, at least in principle, we can model all those algorithms using rule execution.

Even if the computational hypothesis is wrong and cognition is carried out in some noncomputational way, it is still reasonable to use rules in cognitive models when rules provide a descriptive language that is both explanatorily adequate and easy to use. This use of rules is akin to a programmer's use of high-level programming languages like C or Pascal rather than assembly language. At present, rule-based systems should not be removed as a technique for cognitive explanation when all that has been demonstrated so far is that *some* cognitive processes are not well characterized as being rule based and that cognitive science often uses rules that are too coarse grained.

LESSONS FOR COGNITIVE MODELING AND THEORIZING

So far, we have suggested that all cognitive scientists (even antirepresentationalists) accept the importance of mediating states, though they may differ in the particular properties of mediating states that are important for cognitive processing. It is important to note that most of the antirepresentationalist arguments are motivated by key insights about cognitive processing. These insights reflect places where cognitive theorizing has gone astray because of a commitment to a particular representational formalism or the adoption of the wrong model thought. To make these insights more concrete, we draw three lessons from the antirepresentationalists for cognitive modeling (see Table 2). All of these lessons can be adopted while still assuming that cognitive systems have mediating states; that they sometimes contain enduring, discrete symbols, that these symbols can be combined into compositionally

TABLE 2
Three Proposals for the Use of Representations

1. Cognitive models must adopt multiple approaches to representation.
2. Cognitive models must use representations at multiple grainsizes.
3. Cognitive models must be clear about the specification of processes.

complex structures; that some representations are quite abstract; and that, at least in some cases, rules govern some of our behavior.

Cognitive Models Must Adopt Multiple Approaches to Representation

In many arguments about mental representation, for any given proposal for a representational system, an opponent is likely to respond with, ‘‘How can you explain X?’’ where X is some key psychological datum. Some researchers are holding out for the discovery of one representational scheme that explains all of cognitive processing. We suggest that there is no one correct representational scheme. Instead, cognitive science needs multiple approaches to representation. Our argument for this claim is simply that the explananda in psychology and AI (the X’s) range over vastly dissimilar *levels* of processing. No one representational format can handle all of these levels. The complexity of the cognitive system demands that we use diverse representations, methodologies, and styles of explanation.

Our discussion under ‘‘Five Properties That Can Be Added to Mediating States’’ is a clear case in point. We briefly discussed many kinds of representations, but obviously no one of them is *the* representation (except, of course, our construal of representations as mediating states—but this construal works because it is rather weak, although it is far from vacuous). Different kinds of representations are suited for different psychological processes (see Markman, 1999, for an extended discussion of representational assumptions made in psychology). For example, even a transient, mediating state with no discrete symbols is perfectly appropriate as a steam engine governor or a thermostat. Of course, we are not calling these cognitive systems, but there are low-level psychological processes, such as those involved in the execution of some motor functions, that use simple mediating states of this type. Furthermore, the cognitive states underlying many visual events, particularly those that are not attended to, seem to be transient. Grimes (1996) demonstrated that people fail to recognize significant changes in visual displays when these changes fall in peripheral visual areas that are part of the visual field, but are not the focus of attention (see also Simons & Levin, 1998).

Cognitive science should not try to shoehorn cognition and intelligence (human or otherwise) into one model. On our multifaceted view of representation, not only must we stop looking for *the* one true representational scheme, but we must also stop bludgeoning proposals for representational schemes with arguments that they fail to explain psychological processes for which they were never intended. The diversity of representational schemes is to be embraced rather than avoided.

Failure to appreciate the importance of, and need for, representational diversity harms psychology another way. Frequently, when cognitive scientists (philosophers, mainly) find problems with particular representational formats or even paradigms, they throw out not only the representation, but the computational hypothesis too, blaming it as the ultimate culprit rather than just the representation paradigms (see, e.g., Goel, 1995). This reaction seems misguided because *the* major insight in cognitive science on which everything else rests is the Computational Hypothesis (Dietrich, 1994). The Computational Hypothesis in cognitive science is not committed to any particular representational paradigm, so showing that gross rules are *not* explanatorily useful does not prove that computationalism is false.

Arguing for diversity is one thing; actually doing it is something else. It can be quite difficult to generate models that reconcile different representational formats. That is one reason why this approach has not had widespread appeal. However, there are models that use different formats within the same system. For example, Forbus, Gentner, and Law's (1995) MAC/FAC model of analogical retrieval uses two stages, each of which is mediated by a different type of representation. Every item in memory has both a feature representation in which features are not connected to each other and also a representation in which there is role–argument structure. In the first stage, a feature representation of a memory probe is compared to the items in long-term memory to find those with a high degree of overall featural overlap. Then, the structured representations of a small number of memory items with a high degree of overlap take part in a computationally intensive comparison with a structured representation of the probe. The best matching memory items from this structural comparison process are retrieved from memory. In this model, a simple featural representation is used when a computationally cheap process is required as a first filter. A more complex, structured representation is used when the number of memory items has been winnowed enough so that there are sufficient resources for more complex processing. This model adequately captures psychological evidence, demonstrating that judgments of similarity are more sensitive to the structure in mental representations than is the ability to retrieve analogs from memory.

Another example is Anderson's (1983a, 1993) ACT system. This model incorporates both a semantic network and rules. The semantic network has nodes corresponding to concepts and links that connect related concepts. Nodes in the network are activated to the degree they are present in working memory. Cognitive procedures are then carried out by matching the antecedents for rules against the contents of working memory. Thus, this system is also able to coordinate the use of more than one kind of mental representation.

In summary, the cognitive system is unlikely to be based on only one type of representation. There are too many different kinds of processes that vary in how quickly they must be carried out and in their computational complexity for that to be possible. Instead, cognitive science must find ways to integrate processes involving different kinds of representations.

Cognitive Models Must Adopt Representations at Multiple Grainsizes

One issue that arises repeatedly is that many symbolic representations do not capture the fine details of cognition. For example, having a symbol ‘‘cup’’ that is active whenever the concept ‘‘cup’’ is needed fails to account for the context sensitivity of our cup concept(s). This point alone suggests that cognitive models should use microfeatures so that the set of active features is different every time a particular concept is encountered. However, this proposal misses the insight that there is something that is also invariant when we use a concept. One solution to this problem is to assume that cognitive representations for the same item exist at many different grainsizes. Some representational elements may be roughly equivalent to the symbol ‘‘cup,’’ representing cups in general. Other elements may represent only subtle microfeatures of cups or many different kinds of cups. This way, there will be more than one way to represent the same item. This redundancy allows models to deal with cases in which context sensitivity is important and also to deal with cases in which invariance across context is important.

This suggestion of ours about multiple grainsizes is a version of our previous suggestion that many different representational paradigms be used in psychological theorizing. We are saying that even within a single paradigm, representations should accommodate multiple levels information.

As one example of how multiple grainsizes can be incorporated into a single model, we consider Hummel and Holyoak’s (1997) LISA, a connectionist model of analogical reasoning. LISA uses both distributed representations and local representations of concepts. The distributed representation consists of a pool of features for a concept. A concept is active when some set of those features of that concept have a high level of activation. The local representation consists of labeled nodes that correspond to different objects and relations. For example, the concept ‘‘John’’ can be accessed in its general form by activating its node. Activating this node will in turn activate features associated with this concept. The set of features that is activated is different in different contexts. Thus, LISA is able both to respond to the context of a situation and to abstract away from the fine details of a particular situation when necessary. This representational strategy seems appropriate for modeling people’s use of analogies.

The idea that cognitive models need representations at a variety of grainsizes is also consistent with neurobiological evidence of song production in the zebra finch. Behavioral evidence suggests that these bird songs have both syllables (typical patterns of notes) and also specific notes. Yu and Margoliash (1996) have found that activity in the HVc nucleus of the forebrain in zebra finches is correlated with production of song syllables, while activity in the robustus archistriatalis is correlated with production of individual notes of the zebra finch song. These results suggest that there is a hierarchical control for the song in the zebra finch. Thus, the fluent operation of the song system requires representations of different grainsizes within a single system. This general principle is likely to operate over a wide range of psychological processes.

Cognitive Models Must be Clear About the Specification of Processes

We stated that a requirement for something to count as a mediating state was that there be some process that extracts and uses information it contains—the mediating state must *mediate.* One reason why this specification is important is that the implications of a representational formalism for psychological models are not clear until the processing assumptions have been laid out (see also Anderson, 1978, and Palmer, 1978). Another reason is that in an important sense, without specifying the processing assumptions, there is no mediating state at all and hence no representation at all (there is at best only potential for representation), so all further discussions beg many important questions.

For example, Smolensky (1991) discusses a possible context-dependent representation for a cup. He suggests that the microfeatures that are active in the context of a cup with coffee may be very different than the microfeatures that are active in other situations. On the surface, this proposal seems to make clear that a distributed representation can be used as a model of the context dependence of concepts and that this type of representation resolves a problem with traditional symbolic models. However, this proposal does not carry with it any proposals for how the features are activated or any proposals for how such microfeatures would be used by other cognitive processes (or, for that matter, for how the microfeatures manage to represent a cup in the first place). Thus, on its own, it is not clear what such a representation is capable of doing and hence it is less than clear that it can live up to Smolensky's claims for it. Indeed, vector representations like those used in many distributed connectionist models, when combined with the dot product as a mechanism for comparing vector pairs, are unlikely to be sufficient to model the complexity of human similarity processing because of the need to identify specific commonalities and differences arising from comparisons (Gentner & Markman, 1997).

Representations like activation vectors are appealing because they seem to embody the flexibility often observed in cognitive processing. However, a representation is only fluid if there is an associated process that allows it to be used flexibly. No representation is fluid by itself. A dot-product comparison process only allows the computation of proximity between two vectors, thereby limiting a model's flexibility, not increasing it. Conversely, suitable processes applied to highly structured representations can make them very flexible. Indeed, one major appeal of a universal grammar in linguistics is that it allows infinite productivity from a finite number of elements and rules.

Barnden (1994) suggests that the process of analogical reasoning may allow symbol systems to exhibit some of the flexibility often associated with connectionist models. For example, as discussed above, models of analogy and similarity assume that the arguments of relations that have been placed in correspondence are themselves also placed in correspondence (Gentner, 1983, 1989; Gentner & Markman, 1997; Holyoak & Thagard, 1994). This principle of parallel connectivity allows nonidentical items to be matched. Thus, symbol systems are not rigidly required to match only to elements that have identical symbol names. The structure of the representations can allow nonidentical elements to correspond. Further, just as connectionist models exhibit graceful degradation, so too do some symbolic models of analogy and similarity. If information from one domain is missing, it can be filled in by analogy to a second domain by carrying over relations that are connected to structure that already matches.

Carrying this example further, the specification of processing assumptions can also bring about unforeseen flexibility in other ways. Many models of analogical reasoning are able to form multiple interpretations of a single match. For example, given a ‘‘double’’ metaphor like ‘‘A cloud is a sponge,’’ people can generate more than one interpretation (e.g., ‘‘both are fluffy’’ or ‘‘both can hold water’’). Likewise, when given structured representations of this metaphor, models of analogy, like SME (Falkenhainer, et al., 1989) and LISA (Hummel & Holyoak, 1997), can form both interpretations. These models are able to construct both interpretations because they are able to enforce a set of constraints (e.g., structural consistency in the case of analogy) and then start over and form a different interpretation. Interestingly, models that do not strictly enforce constraints, like the ACME model, which uses a process of parallel constraint satisfaction, are unable to form multiple interpretations of a comparison (Holyoak & Thagard, 1989). The ability to form multiple interpretations is also evident in the perception of the Necker cube, which flips between stable three dimensional interpretations.

As this excursion into models of analogical reasoning demonstrates, it is critical to make both representational assumptions *and* processing assumptions explicit in order to understand the implications of a particular representational format for cognitive processing and to avoid begging important questions about when something is a representation and how it succeeds in representing.

CONCLUSION

Arguments against the need for representations typically have the form ‘‘Some cognitive process C does not require internal states with property X, therefore no cognitive processes require internal states with property X. Furthermore, because X is constitutive of representations in general, it follows that representations as such are not needed.’’ The five properties that can be added to mediating states that we examined previously have all been the subject of arguments of this type.

This argument form is a hasty generalization. It is clear for each of the five properties discussed previously that they are not needed to explain certain cognitive processes, but they are required to explain others. The term ‘‘hasty’’ is used advisedly here because a central complaint of anti-representationalists is that cognitive science has not made enough progress in a representational paradigm to warrant continued use of representations. Instead, they argue, cognitive science consists of a collection of nearly independent micro-theories, each supported by a body of data.

The arguments in this article directly address the issue of progress. We suggest that none of the five properties discussed are themselves constitutive of representation, since in each case the property can be removed and yet the central condition on representation remains. That insight was the basis of the definition of a mediating state. Indeed, it appears that there is nothing more to being a representation than being a mediating state. Mediating states not only constitute the general class to which more specific kinds of representations belong; they capture the essence of representation. Instead of debating whether representations exist or what the one true representational formalism is, cognitive science can make more progress by studying which properties of mediating states (i.e., representations) are needed to explain particular classes of cognitive processes. This article is a defense of representation because it suggests that all cognitive scientists accept the core properties of representation. Debates over representation in cognitive science are actually debates about what additional properties of representations are necessary to understand cognitive processing. Where the debate over representation goes awry is in assuming that there is be only *one* set of properties that can suffice for all cognitive processes.

By leaving aside debates over the existence of representations, cognitive science can focus on the crucial issue of what kinds of representations are used by different cognitive processes and how then these representations come to have their content.

REFERENCES

Agre, P. E. (1995). Computational research on interaction and agency. *Artificial Intelligence,* **72,** 1–52.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review,* **85,** 249–277.

Anderson, J. R. (1983a). *The architecture of cognition.* Cambridge, MA: Harvard Univ. Press.

Anderson, J. R. (1983b). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior,* **22,** 261–295.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Baillargeon, R. (1987). Object permanence in 3.5- and 4.5-month old infants. *Developmental Psychology,* **23,** 655–664.

Barnden, J. A. (1994). On the connectionist implementation of analogy and working memory matching. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neutral computation theory: Analogy, metaphor, and reminding (Vol. 3, pp. 327–374). Norwood, NJ: Ablex.*

Barsalou, L. W. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences* **22**(4), 577–660.

Bassok, M., Chase, V. M., & Martin, S. A. (1998). Adding apples and oranges: Semantic constraints on application of formal rules. *Cognitive Psychology,* **35**(2), 99–134.

Bickhard, M. & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science.* Amsterdam: North-Holland/Elsevier.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review,* **94**(2), 115–147.

Brooks, L. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence,* **47,** 139–159.

Chalmers, D. J. (1990). Why Fodor and Pylyshyn were wrong: The simplest refutation. In *The Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 340–347). Cambridge, MA: Lawrence Erlbaum Associates.

Cheng, P. W., & Holyoak, K. J. (1989). On the natural selection of reasoning theories. *Cognition,* **33,** 285–313.

Chomsky, N., & Halle, M. (1991). *The sound pattern of English.* Cambridge, MA: The MIT Press.

Clancey, W. J. (1997). *Situated action: On human knowledge and computer representations.* New York: Cambridge Univ. Press.

Clark, A. (1993). *Associative engines: Connectionism, concepts, and representational change.* Cambridge, MA: The MIT Press.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic priming. *Psychological Review,* **82**(6), 407–428.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition,* **31**(3), 187–276.

Dietrich, E. (1990). Computationalism. *Social Epistemology,* **4**(2), 135–154.

Dietrich, E. (1994). Thinking computers and the problem of intentionality. In E. Dietrich (Ed.) *Thinking computers and virtual persons.* San Diego: Academic Press.

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

Elman, J. L. (1990). Finding structure in time. *Cognitive Science,* **14**(2), 179–212.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence,* **41**(1), 1–63.

Fodor, J., & McLaughlin, B. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition,* **35,** 183–204.

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

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science,* 19, 141–205.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science,* **7,** 155–170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). New York: Cambridge Univ. Press.

Gentner, D., & Markman, A. B. (1997). Structural alignment in analogy and similarity. *American Psychologist,* **52**(1), 45–56.

Goel, V. (1995). *Sketches of thought.* Cambridge, MA: The MIT Press.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception* (pp. 89–110). New York: Oxford Univ. Press.

Herskovits, A. (1986). *Language and spatial cognition: An interdisciplinary study of the prepositions in English.* New York: Cambridge Univ. Press.

Higgins, E.T. (1992). Achieving a 'shared reality' in the communication game: A social action that creates meaning. *Journal of Language and Social Psychology,* **11,** 107–131.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science,* **13**(3), 295–355.

Holyoak, K. J., & Thagard, P. (1994). *Mental leaps.* Cambridge, MA: The MIT Press.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review,* **99**(3), 480–517.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review,* **104**(3), 427–466.

Hutchins, E. (1995). *Cognition in the wild.* Cambridge, MA: The MIT Press.

Ivry, R. B., & Robertson, L. C. (1998). *The two sides of perception.* Cambridge, MA: The MIT Press.

Johnson-Laird, P. N. (1983). *Mental models.* New York: Cambridge Univ. Press.

Johnson-Laird, P. N., Byrne, R. M., & Tabossi, P. (1989). Reasoning by model: The case of multiple quantification. *Psychological Review,* **96**(4), 658–673.

Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction.* Hillsdale, NJ: Erlbaum.

Kaplan, A. S., & Medin, D. L. (1997). The coincidence effect in similarity and choice. *Memory and Cognition,* **25**(4), 570–576.

Keane, M. T., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science,* **18,** 387–438.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* Cambridge, MA: The MIT Press.

Keysar, B., & Bly, B. (1995). Intuitions of the transparency of idioms: Can one keep a secret by spilling the beans? *Journal of Memory and Language,* **34**(1), 89–109.

Kim, J. J., Pinker, S., Prince, A., & Parasada, S. (1991). Why no mere mortal has ever flown out to center field. *Cognitive Science,* **15**(2), 173–218.

Klein, M., Shapiro, E., & Kandel, E. R. (1980). Synaptic plasticity and the modulation of the Ca++ current. *Journal of Experimental Biology,* **89,** 117–157.

Kolodner, J. (1993). *Case-based reasoning.* San Mateo, CA: Morgan Kaufmann.

Koriat, A. (1994). Memory's knowledge of its own knowledge: The accessibility account of the feeling of knowing. In J. Metcalfe & A.P. Shimamura (Eds.), *Metacognition.* Cambridge, MA: The MIT Press.

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

Landau, B., & Jackendoff, R. (1993). ''What'' and ''where'' in spatial language and spatial cognition. *Behavioral and Brain Sciences,* **16**(2), 217–266.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review,* **104**(2), 211–240.

Langacker, R. W. (1986). An introduction to cognitive grammar. *Cognitive Science,* **10**(1), 1–40.

Lewin, K. (1935). *A dynamic theory of personality.* New York: McGraw–Hill.

Marcus, G. F., Brinkmann, U., Clahsen, H., Wiese, R., & Pinker, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology,* **29,** 189–256.

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Erlbaum.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology,* **25**(4), 431–467.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification. *Psychological Review,* **85**(3), 207–238.

Metcalfe, J., & Shimamura, A. P. (Eds.) (1994). *Metacognition: Knowing about knowing.* Cambridge, MA: The MIT Press.

Newell, A. (1990). Unified theories of cognition. Cambridge, MA: Harvard Univ. Press.

Newell, A., & Simon, H. A. (1963). GPS: A program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought.* R. Oldenbourg, Munich.

Norman, D. A., & Rumelhart, D. E. (1975). *Explorations in cognition.* San Francisco: W. H. Freeman.

Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization relationship. *Journal of ExperimentalPsychology: General,* **115**(1), 39–57.

Novick, L. R. (1990). Representational transfer in problem solving. *Psychological Science,* **1**(2), 128–132.

Palmer, S. E. (1977). Hierarchical structure in perceptual representations. *Cognitive Psychology,* **9,** 441–474.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B.B. Lloyd (Eds.) Cognition and categorization (pp. 259–302). Hillsdale, NJ: Erlbaum.

Perner, J. F. (1991). *Understanding the representational mind.* Cambridge, MA: The MIT Press.

Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence,* **46**(1–2), 77–106.

Pollack, J. L. (1994). Justification and defeat. *Artificial Intelligence,* **67,** 377–407.

Reder, L. M. (1996). *Implicit memory and metacognition.* Mahwah, NJ: Erlbaum.

Reder, L. M., & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **18**(3), 435–451.

Reed, S. K., & Bolstad, C. A. (1991). Use of examples and procedures in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **17**(4), 753–766.

Regier, T. (1996). *The human semantic potential.* Cambridge, MA: The MIT Press.

Rips, L. J. (1994). *The psychology of Proof: Deductive reasoning in human thinking.* Cambridge, MA: The MIT Press.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology,* **16,** 371–416.

Schank, R. C., & Abelson, R. (1977). *Scripts, plans, goals and understanding.* Hillsdale, NJ: Erlbaum.

Schank, R. C., Kass, A., & Riesbeck, C. K. (1994). *Inside case-based explanation.* Hillsdale, NJ: Erlbaum.

Shannon, C. (1949). *The mathematical theory of communication.* Urbana, IL: Univ. of Illinois Press.

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review,* **5,** 644–649.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research,* **16,** 158–174.

Smith, E. E., & Medin, D.L. (1981). *Categories and concepts.* Cambridge, MA: Harvard Univ. Press.

Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 146–178) New York: Cambridge Univ. Press.

Smith, L. B., & Sera, M. D. (1992). A developmental analysis of the polar structure of dimensions. *Cognitive Psychology,* **24**(1), 99–142.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences,* **11**(1), 1–74.

Smolensky, P. (1991). Connectionism, constituency, and the language of thought. In B. Loewer & G. Rey (Eds.), Meaning in mind: Fodor and his critics. Cambridge, MA: Blackwell.

Thelen, E. (1995). Time-scale dynamics and the development of an embodied cognition. In R.F. Port & T. van Gelder (Eds.), *Mind as motion* (pp. 69–100). Cambridge, MA: The MIT Press.

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

Tversky, A. (1977). Features of similarity. *Psychological Review,* **84**(4), 327–352.

Tversky, A., & Gati, I. (1982). Similarity, separability and the triangle inequality. *Psychological Review,* **89**(2), 123–154.

Van Gelder, T., & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. In R. F. Port & T. van Gelder(Eds.), *Mind as motion.* Cambridge, MA: The MIT Press.

Vera, A.H., & Simon, H.A. (1993). Situated action: A symbolic interpretation. *Cognitive Science,* **17,** 7–48.

Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of reasoning structure and content.* London: Routledge.

Wellman, H. M. (1990). *The child's theory of mind.* Cambridge, MA: The MIT Press.

Yu, A.C., & Margoliash, D. (1996). Temporal hierarchical control of singing in birds. *Science,* **273,** 1871–1875.

(Accepted August 11, 1999)