

Perceptual symbol systems

Lawrence W. Barsalou

Department of Psychology, Emory University, Atlanta, GA 30322

barsalou@emory.edu userwww.service.emory.edu/~barsalou/

Abstract: Prior to the twentieth century, theories of knowledge were inherently perceptual. Since then, developments in logic, statistics, and programming languages have inspired amodal theories that rest on principles fundamentally different from those underlying perception. In addition, perceptual approaches have become widely viewed as untenable because they are assumed to implement recording systems, not conceptual systems. A perceptual theory of knowledge is developed here in the context of current cognitive science and neuroscience. During perceptual experience, association areas in the brain capture bottom-up patterns of activation in sensory-motor areas. Later, in a top-down manner, association areas partially reactivate sensory-motor areas to implement perceptual symbols. The storage and reactivation of perceptual symbols operates at the level of perceptual components – not at the level of holistic perceptual experiences. Through the use of selective attention, schematic representations of perceptual components are extracted from experience and stored in memory (e.g., individual memories of *green*, *purr*, *hot*). As memories of the same component become organized around a common frame, they implement a simulator that produces limitless simulations of the component (e.g., simulations of *purr*). Not only do such simulators develop for aspects of sensory experience, they also develop for aspects of proprioception (e.g., *lift*, *run*) and introspection (e.g., *compare*, *memory*, *happy*, *hungry*). Once established, these simulators implement a basic conceptual system that represents types, supports categorization, and produces categorical inferences. These simulators further support productivity, propositions, and abstract concepts, thereby implementing a fully functional conceptual system. Productivity results from integrating simulators combinatorially and recursively to produce complex simulations. Propositions result from binding simulators to perceived individuals to represent type-token relations. Abstract concepts are grounded in complex simulations of combined physical and introspective events. Thus, a perceptual theory of knowledge can implement a fully functional conceptual system while avoiding problems associated with amodal symbol systems. Implications for cognition, neuroscience, evolution, development, and artificial intelligence are explored.

Keywords: analogue processing; categories; concepts; frames; imagery; images; knowledge; perception; representation; sensory-motor representations; simulation; symbol grounding; symbol systems

The habit of abstract pursuits makes learned men much inferior to the average in the power of visualization, and much more exclusively occupied with words in their “thinking.”

Bertrand Russell (1919b)

1. Introduction

For the last several decades, the fields of cognition and perception have diverged. Researchers in these two areas know ever less about each other’s work, and their discoveries have had diminishing influence on each other. In many universities, researchers in these two areas are in different programs, and sometimes in different departments, buildings, and university divisions. One might conclude from this lack of contact that perception and cognition reflect independent or modular systems in the brain. Perceptual systems pick up information from the environment and pass it on to separate systems that support the various cognitive functions, such as language, memory, and thought. I will argue that this view is fundamentally wrong. Instead, cognition is inherently perceptual, sharing systems with perception at both the cognitive and the neural levels. I will further suggest that the divergence between cognition and perception reflects the widespread assumption that cognitive representations are inherently nonperceptual, or what I will call *amodal*.

1.1. Grounding cognition in perception

In contrast to modern views, it is relatively straightforward to imagine how cognition could be inherently perceptual.

As Figure 1 illustrates, this view begins by assuming that perceptual states arise in sensory-motor systems. As discussed in more detail later (sect. 2.1), a perceptual state can contain two components: an unconscious neural representation of physical input, and an optional conscious experience. Once a perceptual state arises, a subset of it is extracted via selective attention and stored permanently in



LAWRENCE BARSALOU is Professor of Psychology at Emory University, having previously held positions at the Georgia Institute of Technology and the University of Chicago. He is the author of over 50 scientific publications in cognitive psychology and cognitive science. His research addresses the nature of human knowledge and its roles in language, memory, and thought,

focusing specifically on the perceptual bases of knowledge, the structural properties of concepts, the dynamic representation of concepts, the construction of categories to achieve goals, and the situated character of conceptualizations. He is a Fellow of the American Psychological Society and of the American Psychological Association, he served as Associate Editor for the *Journal of Experimental Psychology: Learning, Memory, and Cognition*, and he is currently on the governing board of the Cognitive Science Society.




Figure 1. The basic assumption underlying perceptual symbol systems: Subsets of perceptual states in sensory-motor systems are extracted and stored in long-term memory to function as symbols. As a result, the internal structure of these symbols is modal, and they are analogically related to the perceptual states that produced them.

long-term memory. On later retrievals, this perceptual memory can function symbolically, standing for referents in the world, and entering into symbol manipulation. As collections of perceptual symbols develop, they constitute the representations that underlie cognition.

Perceptual symbols are modal and analogical. They are modal because they are represented in the same systems as the perceptual states that produced them. The neural systems that represent color in perception, for example, also represent the colors of objects in perceptual symbols, at least to a significant extent. On this view, a common representational system underlies perception and cognition, not independent systems. Because perceptual symbols are modal, they are also analogical. The structure of a perceptual symbol corresponds, at least somewhat, to the perceptual state that produced it.¹

Given how reasonable this perceptually based view of cognition might seem, why has it not enjoyed widespread acceptance? Why is it not in serious contention as a theory of representation? Actually, this view dominated theories of mind for most of recorded history. For more than 2,000 years, theorists viewed higher cognition as inherently perceptual. Since Aristotle (4th century BC/1961) and Epicurus (4th century BC/1994), theorists saw the representations that underlie cognition as imagistic. British empiricists such as Locke (1690/1959), Berkeley (1710/1982), and Hume (1739/1978) certainly viewed cognition in this manner. Images likewise played a central role in the theories of later nativists such as Kant (1787/1965) and Reid (1764/1970; 1785/1969). Even recent philosophers such as Russell (1919b) and Price (1953) have incorporated images centrally into their theories. Until the early twentieth century, nearly all theorists assumed that knowledge had a strong perceptual character.

After being widely accepted for two millennia, this view withered with mentalism in the early twentieth century. At that time, behaviorists and ordinary language philosophers successfully banished mental states from consideration in much of the scientific community, arguing that they were unscientific and led to confused views of human nature (e.g., Ryle 1949; Watson 1913; Wittgenstein 1953). Because perceptual theories of mind had dominated mentalism to that point, attacks on mentalism often included a critique

of images. The goal of these attacks was not to exclude images from mentalism, however, but to eliminate mentalism altogether. As a result, image-based theories of cognition disappeared with theories of cognition.

1.2. Amodal symbol systems

Following the cognitive revolution in the mid-twentieth century, theorists developed radically new approaches to representation. In contrast to pre-twentieth century thinking, modern cognitive scientists began working with representational schemes that were inherently nonperceptual. To a large extent, this shift reflected major developments outside cognitive science in logic, statistics, and computer science. Formalisms such as predicate calculus, probability theory, and programming languages became widely known and inspired technical developments everywhere. In cognitive science, they inspired many new representational languages, most of which are still in widespread use today (e.g., feature lists, frames, schemata, semantic nets, procedural semantics, production systems, connectionism).

These new representational schemes differed from earlier ones in their relation to perception. Whereas earlier schemes assumed that cognitive representations utilize perceptual representations (Fig. 1), the newer schemes assumed that cognitive and perceptual representations constitute separate systems that work according to different principles. Figure 2 illustrates this assumption. As in the framework for perceptual symbol systems in Figure 1, perceptual states arise in sensory-motor systems. However, the next step differs critically. Rather than extracting a subset of a perceptual state and storing it for later use as a symbol, an amodal symbol system transduces a subset of a perceptual state into a completely new representation language that is inherently nonperceptual.

As amodal symbols become transduced from perceptual states, they enter into larger representational structures, such as feature lists, frames, schemata, semantic networks, and production systems. These structures in turn constitute a fully functional symbolic system with a combinatorial syntax and semantics, which supports all of the higher cognitive functions, including memory, knowledge, language, and thought. For general treatments of this approach, see Dennett (1969), Newell and Simon (1972), Fodor (1975), Pylyshyn (1984), and Haugeland (1985). For reviews of specific theories in psychology, see E. Smith and Medin (1981), Rumelhart and Norman (1988), and Barsalou and Hale (1993).

It is essential to see that the symbols in these systems are amodal and arbitrary. They are amodal because their internal structures bear no correspondence to the perceptual states that produced them. The amodal symbols that represent the colors of objects in their absence reside in a different neural system from the representations of these colors during perception itself. In addition, these two systems use different representational schemes and operate according to different principles.

Because the symbols in these symbol systems are amodal, they are linked arbitrarily to the perceptual states that produce them. Similarly to how words typically have arbitrary relations to entities in the world, amodal symbols have arbitrary relations to perceptual states. Just as the word "chair" has no systematic similarity to physical chairs, the amodal symbol for *chair* has no systematic similarity to




Figure 2. The basic assumption underlying amodal symbol systems: Perceptual states are transduced into a completely new representational system that describes these states amodally. As a result, the internal structure of these symbols is unrelated to the perceptual states that produced them, with conventional associations establishing reference instead.

perceived chairs. As a consequence, similarities between amodal symbols are not related systematically to similarities between their perceptual states, which is again analogous to how similarities between words are not related systematically to similarities between their referents. Just as the words “blue” and “green” are not necessarily more similar than the words “blue” and “red,” the amodal symbols for *blue* and *green* are not necessarily more similar than the amodal symbols for *blue* and *red*.²

Amodal symbols bear an important relation to words and language. Theorists typically use linguistic forms to represent amodal symbols. In feature lists, words represent features, as in:

CHAIR (1)
seat
back
legs

Similarly in schemata, frames, and predicate calculus expressions, words represent relations, arguments, and values, as in:

EAT (2)
Agent = horse
Object = hay

Although theorists generally assume that words do not literally constitute the content of these representations, it is assumed that close amodal counterparts of words do. Although the word “horse” does not represent the value of *Agent* for *EAT* in (2), an amodal symbol that closely parallels this word does. Thus, symbolic thought is assumed to be analogous in many important ways to language. Just as language processing involves the sequential processing of words in a sentence, so conceptual processing is assumed to involve the sequential processing of amodal symbols in list-like or sentence-like structures (e.g., Fodor & Pylyshyn 1988).

It is important to see that this emphasis on amodal and arbitrary symbols also exists in some, but not all, connectionist schemes for representing knowledge (e.g., McClelland et al. 1986; Rumelhart et al. 1986). Consider a feed-forward network with back propagation. The input units in the first layer constitute a simple perceptual system that

codes the perceived features of presented entities. In contrast, the internal layer of hidden units is often interpreted as a simple conceptual system, with a pattern of activation providing the conceptual representation of an input pattern. Most importantly, the relation between a conceptual representation and its perceptual input is arbitrary for technical reasons. Prior to learning, the starting weights on the connections between the input units and the hidden units are set to small *random* values (if the values were all 0, the system could not learn). As a result, the conceptual representations that develop through learning are related arbitrarily to the perceptual states that activate them. With different starting weights, arbitrarily different conceptual states correspond to the same perceptual states. Even though connectionist schemes for representation differ in important ways from more traditional schemes, they often share the critical assumption that cognitive representations are amodal and arbitrary.

Connectionist representational schemes need not necessarily work this way. If the same associative network represents information in both perception and cognition, it grounds knowledge in perception and is not amodal (e.g., Pulvermüller 1999). As described later (sects. 2.2.1, 2.5), shared associative networks provide a natural way to view the representation of perceptual symbols.

1.2.1. Strengths. Amodal symbol systems have many powerful and important properties that any fully functional conceptual system must exhibit. These include the ability to represent types and tokens, to produce categorical inferences, to combine symbols productively, to represent propositions, and to represent abstract concepts. Amodal symbol systems have played the critical role of making these properties central to theories of human cognition, making it clear that any viable theory must account for them.

1.2.2. Problems. It has been less widely acknowledged that amodal symbol systems face many unresolved problems. First, there is little direct empirical evidence that amodal symbols exist. Using picture and word processing tasks, some researchers have explicitly tested the hypothesis that conceptual symbols are amodal (e.g., Snodgrass 1984; Theios & Amrhein 1989). However, a comprehensive review of this work concluded that conceptual symbols have a perceptual character (Glaser 1992; also see Seifert 1997). More recently, researchers have suggested that amodal vectors derived from linguistic context underlie semantic processing (Burgess & Lund 1997; Landauer & Dumais 1997). However, Glenberg et al. (1998b) provide strong evidence against these views, suggesting instead that affordances derived from sensory-motor simulations are essential to semantic processing.

Findings from neuroscience also challenge amodal symbols. Much research has established that categorical knowledge is grounded in sensory-motor regions of the brain (for reviews see Damasio 1989; Gainotti et al. 1995; Pulvermüller 1999; also see sect. 2.3). Damage to a particular sensory-motor region disrupts the conceptual processing of categories that use this region to perceive physical exemplars. For example, damage to the visual system disrupts the conceptual processing of categories whose exemplars are primarily processed visually, such as *birds*. These findings strongly suggest that categorical knowledge is not amodal.³

In general, the primary evidence for amodal symbols is indirect. Because amodal symbols can implement conceptual systems, they receive indirect support through their instrumental roles in these accounts. Notably, however, amodal symbols have not fared well in implementing all computational functions. In particular, they have encountered difficulty in representing spatio-temporal knowledge, because the computational systems that result are cumbersome, brittle, and intractable (e.g., Clark 1997; Glasgow 1993; McDermott 1987; Winograd & Flores 1987). Although amodal symbol systems do implement some computational functions elegantly and naturally, their inadequacies in implementing others are not encouraging.

Another shortcoming of amodal symbol systems is their failure to provide a satisfactory account of the transduction process that maps perceptual states into amodal symbols (Fig. 2). The lack of an account for such a critical process should give one pause in adopting this general framework. If we cannot explain how these symbols arise in the cognitive system, why should we be confident that they exist? Perhaps even more serious is the complete lack of cognitive and neural evidence that such a transduction process actually exists in the brain.

A related shortcoming is the symbol grounding problem (Harnad 1987; 1990; Newton 1996; Searle 1980), which is the converse of the transduction problem. Just as we have no account of how perceptual states become mapped to amodal symbols during transduction, neither do we have an account of how amodal symbols become mapped back to perceptual states and entities in the world. Although amodal theories often stress the importance of symbol interpretation, they fail to provide compelling accounts of the interpretive scheme that guides reference. Without such an account, we should again have misgivings about the viability of this approach.⁴

A related problem concerns how an amodal system implements comprehension in the absence of physical referents. Imagine that amodal symbols are manipulated to envision a future event. If nothing in the perceived environment grounds these symbols, how does the system understand its reasoning? Because the processing of amodal symbols is usually assumed to be entirely syntactic (based on form and not meaning), how could such a system have any sense of what its computations are about? It is often argued that amodal symbols acquire meaning from associated symbols, but without ultimately grounding terminal symbols, the problem remains unsolved. Certainly people have the experience of comprehension in such situations.

One solution is to postulate mediating perceptual representations (e.g., Harnad 1987; Höffding 1891; Neisser 1967). According to this account, every amodal symbol is associated with corresponding perceptual states in long-term memory. For example, the amodal symbol for *dog* is associated with perceptual memories of dogs. During transduction, the perception of a dog activates these perceptual memories, which activate the amodal symbol for *dog*. During symbol grounding, the activation of the amodal symbol in turn activates associated perceptual memories, which ground comprehension. Problematically, though, perceptual memories are doing all of the work, and the amodal symbols are redundant. Why couldn't the system simply use its perceptual representations of dogs alone to represent *dog*, both during categorization and reasoning?

The obvious response from the amodal perspective is that amodal symbols perform additional work that these perceptual representations cannot perform. As we shall see, however, perceptual representations can play the critical symbolic functions that amodal symbols play in traditional systems, so that amodal symbols become redundant. If we have no direct evidence for amodal symbols, as noted earlier, then why postulate them?

Finally, amodal symbol systems are too powerful. They can explain any finding post hoc (Anderson 1978), but often without providing much illumination. Besides being unfalsifiable, these systems often fail to make strong a priori predictions about cognitive phenomena, especially those of a perceptual nature. For example, amodal theories do not naturally predict distance and orientation effects in scanning and rotation (Finke 1989; Kosslyn 1980), although they can explain them post hoc. Such accounts are not particularly impressive, though, because they are unconstrained and offer little insight into the phenomena.

1.2.3. Theory evaluation. Much has been made about the ability of amodal theories to explain any imagery phenomenon (e.g., Anderson 1978). However, this ability must be put into perspective. If perceptual theories predict these effects a priori, whereas amodal theories explain them post hoc, why should this be viewed as a tie? From the perspective of inferential statistics, Bayesian reasoning, and philosophy of science, post hoc accounts should be viewed with great caution. If a priori prediction is favored over post hoc prediction in these other areas, why should it not be favored here? Clearly, greater credence must be given to a theory whose falsifiable, a priori predictions are supported than to a theory that does not predict these findings a priori, and that accounts for them post hoc only because of its unfalsifiable explanatory power.

Furthermore, the assessment of scientific theories depends on many other factors besides the ability to fit data. As philosophers of science often note, theories must also be evaluated on falsifiability, parsimony, the ability to produce provocative hypotheses that push a science forward, the existence of direct evidence for their constructs, freedom from conceptual problems in their apparatus, and integrability with theory in neighboring fields. As we have seen, amodal theories suffer problems in all these regards. They are unfalsifiable, they are not parsimonious, they lack direct support, they suffer conceptual problems such as transduction and symbol grounding, and it is not clear how to integrate them with theory in neighboring fields, such as perception and neuroscience. For all of these reasons, we should view amodal theories with caution and skepticism, and we should be open to alternatives.

1.3. The current status of perceptual symbol systems

The reemergence of cognition in the mid-twentieth century did not bring a reemergence of perceptually based cognition. As we have seen, representational schemes moved in a nonperceptual direction. Furthermore, theorists were initially hostile to imbuing modern cognitive theories with any perceptual character whatsoever. When Shepard and Metzler (1971) offered initial evidence for image-like representations in working memory (not long-term memory!), they encountered considerable resistance (e.g., Anderson 1978; Pylyshyn 1973; 1981). [See also Pylyshyn: "Computa-

tional Models and Empirical Constraints" *BBS* 1(1) 1978; "Computation and Cognition" *BBS* 3(1) 1980; "Is Vision Continuous with Cognition?" *BBS* 22(3) 1999.] When Kosslyn (1980) presented his theory of imagery, he argued adamantly that permanent representations in long-term memory are amodal, with perceptual images existing only temporarily in working memory (see also Kosslyn 1976).

The reasons for this resistance are not entirely clear. One factor could be lingering paranoia arising from the attacks of behaviorists and ordinary language philosophers. Another factor could be more recent criticisms of imagery in philosophy, some of which will be addressed later (e.g., Dennett 1969; Fodor 1975; Geach 1957). Perhaps the most serious factor has been uncharitable characterizations of perceptual cognition that fail to appreciate its potential. Critics often base their attacks on weak formulations of the perceptual approach and underestimate earlier theorists. As a result, perceptual theories of knowledge are widely misunderstood.

Consider some of the more common misunderstandings: perceptual theories of knowledge are generally believed to contain holistic representations instead of componential representations that exhibit productivity. These theories are widely believed to contain only conscious mental images, not unconscious representations. The representations in these theories are often assumed to arise only in the sensory modalities, not in other modalities of experience, such as proprioception and introspection. These theories are typically viewed as containing only static representations, not dynamic ones. These theories are generally construed as failing to support the propositions that underlie description and interpretation. And these theories are often assumed to include only empirical collections of sense data, not genetically constrained mechanisms.

Careful readings of earlier thinkers, however, indicate that perceptual theories of knowledge often go considerably beyond this simplistic stereotype. Many philosophers, for example, have assumed that perceptual representations are componential and produce representations productively (e.g., Locke, Russell, Price). Many have assumed that unconscious representations, then referred to as "dispositions" and "schemata," produce conscious images (e.g., Locke, Kant, Reid, Price). Many have assumed that images can reflect nonsensory experience, most importantly introspection or "reflection" (e.g., Locke, Hume, Kant, Reid). Many have assumed that images can support the type-token mappings that underlie propositions (e.g., Locke, Reid, Russell, Price). Many have assumed that native mechanisms interpret and organize images (e.g., Kant, Reid). All have assumed that images can be dynamic, not just static, representing events as well as snapshots of time.

As these examples suggest, perceptual theories of knowledge should be judged on the basis of their strongest members, not their weakest. My intent here is to develop a powerful theory of perceptual symbols in the contexts of cognitive science and neuroscience. As we shall see, this type of theory can exhibit the strengths of amodal symbol systems while avoiding their problems.

More and more researchers are developing perceptual theories of cognition. In linguistics, cognitive linguists have made the perceptual character of knowledge a central assumption of their theories (e.g., Fauconnier 1985; 1997; Jackendoff 1987; Johnson 1987; Lakoff 1987; 1988; Lakoff & Johnson 1980; Lakoff & Turner 1989; Langacker 1986;

1987; 1991; 1997; Sweetser 1990; Talmy 1983; 1988; Turner 1996). In psychology, these researchers include Paivio (1971; 1986), Miller and Johnson-Laird (1976), Huttenlocher (1973; 1976), Shannon (1987), J. Mandler (1992), Tomasello (1992), L. Smith (L. Smith & Heise 1992; L. Smith et al. 1992; Jones & L. Smith 1993), Gibbs (1994), Glenberg (1997), Goldstone (Goldstone 1994; Goldstone & Barsalou 1998), Wu (1995), Solomon (1997), MacWhinney (1998), and myself (Barsalou 1993; Barsalou & Prinz 1997; Barsalou et al. 1993; in press). In philosophy, these researchers include Barwise and Etchemendy (1990; 1991), Nersessian (1992), Peacocke (1992), Thagard (1992), Davies and Stone (1995), Heal (1996), Newton (1996), Clark (1997), and Prinz (1997; Prinz & Barsalou, in press a). In artificial intelligence, Glasgow (1993) has shown that perceptual representations can increase computational power substantially, and other researchers are grounding machine symbols in sensory-motor events (e.g., Bailey et al. 1997; Cohen et al. 1997; Rosenstein & Cohen 1998). Many additional researchers have considered the role of perceptual representations in imagery (e.g., Farah 1995; Finke 1989; Kosslyn 1994; Shepard & Cooper 1982; Tye 1991), but the focus here is on perceptual representations in long-term knowledge.

1.4. Recording systems versus conceptual systems

It is widely believed that perceptually based theories of knowledge do not have sufficient expressive power to implement a fully functional conceptual system. As described earlier (sect. 1.2.1), a fully functional conceptual system represents both types and tokens, it produces categorical inferences, it combines symbols productively to produce limitless conceptual structures, it produces propositions by binding types to tokens, and it represents abstract concepts. The primary purpose of this target article is to demonstrate that perceptual symbol systems can implement these functions naturally and powerfully.

The distinction between a recording system and a conceptual system is central to this task (Dretske 1995; Haugeland 1991). Perceptually based theories of knowledge are typically construed as recording systems. A recording system captures physical information by creating attenuated (not exact) copies of it, as exemplified by photographs, videotapes, and audiotapes. Notably, a recording system does not interpret what each part of a recording contains – it simply creates an attenuated copy. For example, a photo of a picnic simply records the light present at each point in the scene without interpreting the types of entities present.

In contrast, a conceptual system interprets the entities in a recording. In perceiving a picnic, the human conceptual system might construe perceived individuals as instances of *tree*, *table*, *watermelon*, *eat*, *above*, and so forth. To accomplish this, the conceptual system binds specific tokens in perception (i.e., individuals) to knowledge for general types of things in memory (i.e., concepts). Clearly, a system that only records perceptual experience cannot construe individuals in this manner – it only records them in the holistic context of an undifferentiated event.

A conceptual system has other properties as well. First, it is inferential, allowing the cognitive system to go beyond perceptual input. Theorists have argued for years that the primary purpose of concepts is to provide categorical inferences about perceived individuals. Again, this is not some-

thing that recording systems accomplish. How does a photo of a dog go beyond what it records to provide inferences about the individual present? Second, a conceptual system is productive in the sense of being able to construct complex concepts from simpler ones. This, too, is not something possible with recording systems. How could a photo of some snow combine with a photo of a ball to form the concept of *snowball*? Third, a conceptual system supports the formulation of propositions, where a proposition results from binding a concept (type) to an individual (token) in a manner that is true or false. Again, this is something that lies beyond the power of recording systems. How does a photo of a dog implement a binding between a concept and an individual?

As long as perceptually based theories of knowledge are viewed as recording systems, they will never be plausible, much less competitive. To become plausible and competitive, a perceptually based theory of knowledge must exhibit the properties of a conceptual system. The primary purpose of this target article is to demonstrate that this is possible.

Of course, it is important to provide empirical support for such a theory as well. Various sources of empirical evidence will be offered throughout the paper, especially in section 4, and further reports of empirical support are forthcoming (Barsalou et al., in press; Solomon & Barsalou 1999a; 1999b; Wu & Barsalou 1999). However, the primary support here will be of a theoretical nature. Because so few theorists currently believe that a perceptually based theory of knowledge could possibly have the requisite theoretical properties, it is essential to demonstrate that it can. Once this has been established, an empirical case can follow.

1.5. Overview

The remainder of this paper presents a theory of perceptual symbols. Section 2 presents six core properties that implement a basic conceptual system: perceptual symbols are neural representations in sensory-motor areas of the brain (sect. 2.1); they represent schematic components of perceptual experience, not entire holistic experiences (sect. 2.2); they are multimodal, arising across the sensory modalities, proprioception, and introspection (sect. 2.3). Related perceptual symbols become integrated into a simulator that produces limitless simulations of a perceptual component (e.g., *red*, *lift*, *hungry*, sect. 2.4). Frames organize the perceptual symbols within a simulator (sect. 2.5), and words associated with simulators provide linguistic control over the construction of simulations (sect. 2.6).

Section 3 presents four further properties, derived from the six core properties, that implement a fully functional conceptual system: simulators can be combined combinatorially and recursively to implement productivity (sect. 3.1); they can become bound to perceived individuals to implement propositions (sect. 3.2). Because perceptual symbols reside in sensory-motor systems, they implement variable embodiment, not functionalism (sect. 3.3). Using complex simulations of combined physical and introspective events, perceptual symbol systems represent abstract concepts (sect. 3.4).

Section 4 sketches implications of this approach. Viewing knowledge as grounded in sensory-motor areas changes how we think about basic cognitive processes, including categorization, concepts, attention, working memory, long-term memory, language, problem solving, decision making,

skill, reasoning, and formal symbol manipulation (sect. 4.1). This approach also has implications for evolution and development (sect. 4.2), neuroscience (sect. 4.3), and artificial intelligence (sect. 4.4).

2. Core properties

The properties of this theory will not be characterized formally, nor will they be grounded in specific neural mechanisms. Instead, this formulation of the theory should be viewed as a high-level functional account of how the brain could implement a conceptual system using sensory-motor mechanisms. Once the possibility of such an account has been established, later work can develop computational implementations and ground them more precisely in neural systems.

Because this target article focuses on the high level architecture of perceptual symbol systems, it leaves many details unspecified. The theory does not specify the features of perception, or why attention focuses on some features but not others. The theory does not address how the cognitive system divides the world into categories, or how abstraction processes establish categorical knowledge. The theory does not explain how the fit between one representation and another is computed, or how constraints control the combination of concepts. Notably, these issues remain largely unresolved in *all* theories of knowledge – not just perceptual symbol systems – thereby constituting some of the field's significant challenges. To provide these missing aspects of the theory would exceed the scope of this article, both in length and ambition. Instead, the goal is to formulate the high-level architecture of perceptual symbol systems, which may well provide leverage in resolving these other issues. From here on, footnotes indicate critical aspects of the theory that remain to be developed.

Finally, this target article proposes a theory of knowledge, not a theory of perception. Although the theory relies heavily on perception, it remains largely agnostic about the nature of perceptual mechanisms. Instead, the critical claim is that whatever mechanisms happen to underlie perception, an important subset will underlie knowledge as well.

2.1. Neural representations in sensory-motor systems

Perceptual symbols are *not* like physical pictures; nor are they mental images or any other form of conscious subjective experience. As natural and traditional as it is to think of perceptual symbols in these ways, this is not the form they take here. Instead, they are records of the neural states that underlie perception. During perception, systems of neurons in sensory-motor regions of the brain capture information about perceived events in the environment and in the body. At this level of perceptual analysis, the information represented is relatively qualitative and functional (e.g., the presence or absence of edges, vertices, colors, spatial relations, movements, pain, heat). The neuroscience literature on sensory-motor systems is replete with accounts of this neural architecture (e.g., Bear et al. 1996; Gazzaniga et al. 1998; Zeki 1993). There is little doubt that the brain uses active configurations of neurons to represent the properties of perceived entities and events.

This basic premise of modern perceptual theory under-

lies the present theory of perceptual symbol systems: a perceptual symbol is a record of the neural activation that arises during perception. Essentially the same assumption also underlies much current work in imagery: common neural systems underlie imagery and perception (e.g., Crammond 1997; Deschaumes-Molinari et al. 1992; Farah 1995; Jeannerod 1994; 1995; Kosslyn 1994; Zatorre et al. 1996). The proposal here is stronger, however, further assuming that the neural systems common to imagery and perception underlie conceptual knowledge as well.

This claim by no means implies that identical systems underlie perception, imagery, and knowledge. Obviously, they must differ in important ways. For example, Damasio (1989) suggests that convergence zones integrate information in sensory-motor maps to represent knowledge. More generally, associative areas throughout the brain appear to play this integrative role (Squire et al. 1993). Although mechanisms outside sensory-motor systems enter into conceptual knowledge, perceptual symbols always remain grounded in these systems. Complete transductions never occur whereby amodal representations that lie in associative areas totally replace modal representations. Thus, Damasio (1989) states that convergence zones “are uninformed as to the content of the representations they assist in attempting to reconstruct. The role of convergence zones is to enact formulas for the reconstitution of fragment-based momentary representations of entities or events in sensory and motor cortices” (p. 46).⁵

2.1.1. Conscious versus unconscious processing. Although neural representations define perceptual symbols, they may produce conscious counterparts on some occasions. On other occasions, however, perceptual symbols function unconsciously, as during preconscious processing and automatized skills. Most importantly, the basic definition of perceptual symbols resides at the neural level: unconscious neural representations – not conscious mental images – constitute the core content of perceptual symbols.⁶

Both the cognitive and neuroscience literatures support this distinction between unconscious neural representations and optional conscious counterparts. In the cognitive literature, research on preconscious processing indicates that conscious states may not accompany unconscious processing, and that if they do, they follow it (e.g., Marcel 1983a; 1983b; Velmans 1991). Similarly, research on skill acquisition has found that conscious awareness falls away as automaticity develops during skill acquisition, leaving unconscious mechanisms largely in control (e.g., Schneider & Shiffrin 1977; Shiffrin 1988; Shiffrin & Schneider 1977). Researchers have similarly found that conscious experience often fails to reflect the unconscious mechanisms controlling behavior (e.g., Nisbett & Wilson 1977). In the neuroscience literature, research on blindsight indicates that unconscious processing can occur in the absence of conscious visual images (e.g., Cowey & Stoerig 1991; Weiskrantz 1986; see also Campion, Lotto & Smith: “Is Blindsight an Effect of Scattered Light, Spared Cortex, and Near-Threshold Vision?” *BBS* 2(3) 1983). Similarly, conscious states typically follow unconscious states when processing sensations and initiating actions, rather than preceding them (Dennett & Kinsbourne 1992; Libet 1982; 1985). Furthermore, different neural mechanisms appear responsible for producing conscious and unconscious processing

(e.g., Farah & Feinberg 1997; Gazzaniga 1988; Schacter et al. 1988).

Some individuals experience little or no imagery. By distinguishing unconscious perceptual processing from conscious perceptual experience, we can view such individuals as people whose unconscious perceptual processing underlies cognition with little conscious awareness. If human knowledge is inherently perceptual, there is no a priori reason it must be represented consciously.

2.2. Schematic perceptual symbols

A perceptual symbol is *not* the record of the entire brain state that underlies a perception. Instead, it is only a very small subset that represents a coherent aspect of the state. This is an assumption of many older theories (e.g., Locke 1690/1959), as well as many current ones (e.g., Langacker 1986; J. Mandler 1992; Talmy 1983). Rather than containing an entire holistic representation of a perceptual brain state, a perceptual symbol contains only a schematic aspect.

The schematic nature of perceptual symbols falls out naturally from two attentional assumptions that are nearly axiomatic in cognitive psychology: Selective attention (1) isolates information in perception, and (2) stores the isolated information in long-term memory. First, consider the role of selective attention in isolating features. During a perceptual experience, the cognitive system can focus attention on a meaningful, coherent aspect of perception. On perceiving an array of objects, attention can focus on the shape of one object, filtering out its color, texture, and position, as well as the surrounding objects. From decades of work on attention, we know that people have a sophisticated and flexible ability to focus attention on features (e.g., Norman 1976; Shiffrin 1988; Treisman 1969), as well as on the relations between features (e.g., Treisman 1993). Although nonselected information may not be filtered out completely, there is no doubt that it is filtered to a significant extent (e.g., Garner 1974; 1978; Melara & Marks 1990).⁷

Once an aspect of perception has been selected, it has a very high likelihood of being stored in long-term memory. On selecting the shape of an object, attention stores information about it. From decades of work on episodic memory, it is clear that where selective attention goes, long-term storage follows, at least to a substantial extent (e.g., Barsalou 1995; F. Craik & Lockhart 1972; Morris et al. 1977; D. Nelson et al. 1979). Research on the acquisition of automaticity likewise shows that selective attention controls storage (Compton 1995; Lassaline & Logan 1993; Logan & Etherton 1994; Logan et al. 1996). Although some nonselected information may be stored, there is no doubt that it is stored to a much lesser extent than selected information. Because selective attention focuses constantly on aspects of experience in this manner, large numbers of schematic representations become stored in memory. As we shall see later, these representations can serve basic symbolic functions. Section 3.1 demonstrates that these representations combine productively to implement compositionality, and section 3.2 demonstrates that they acquire semantic interpretations through the construction of propositions. The use of “perceptual symbols” to this point anticipates these later developments of the theory.

Finally, this symbol formation process should be viewed in terms of the neural representations described in section

2.1. If a configuration of active neurons underlies a perceptual state, selective attention operates on this neural representation, isolating a subset of active neurons. If selective attention focuses on an object's shape, the neurons representing this shape are selected, and a record of their activation is stored. Such storage could reflect the Hebbian strengthening of connections between active neurons (e.g., Pulvermüller 1999), or the indirect integration of active neurons via an adjacent associative area (e.g., Damasio 1989). Conscious experience may accompany the symbol formation process and may be necessary for this process to occur initially, falling away only as a symbol's processing becomes automatized with practice. Most fundamentally, however, the symbol formation process selects and stores a subset of the active neurons in a perceptual state.

2.2.1. Perceptual symbols are dynamic, not discrete. Once a perceptual symbol is stored, it does not function rigidly as a discrete symbol. Because a perceptual symbol is an associative pattern of neurons, its subsequent activation has dynamical properties. Rather than being reinstated exactly on later occasions, its activations may vary widely. The subsequent storage of additional perceptual symbols in the same association area may alter connections in the original pattern, causing subsequent activations to differ. Different contexts may distort activations of the original pattern, as connections from contextual features bias activation toward some features in the pattern more than others. In these respects, a perceptual symbol is an attractor in a connectionist network. As the network changes over time, the attractor changes. As the context varies, activation of the attractor covaries. Thus, a perceptual symbol is neither rigid nor discrete.

2.2.2. Perceptual symbols are componential, not holistic. Theorists often view perceptual representations as conscious holistic images. This leads to various misunderstandings about perceptual theories of knowledge. One is that it becomes difficult to see how a perceptual representation could be componential. How can one construct a schematic image of a shape without orientation combined holistically? If one imagines a triangle consciously, is orientation not intrinsically required in a holistic image?

It may be true that conscious images must contain certain conjunctions of dimensions. Indeed, it may be difficult or impossible to construct a conscious image that breaks apart certain dimensions, such as shape and orientation. If a perceptual symbol is defined as an unconscious neural representation, however, this is not a problem. The neurons for a particular shape could be active, while no neurons for a particular orientation are. During the unconscious processing of perceptual symbols, the perceptual symbol for a particular shape could represent the shape componentially, while perceptual symbols for other dimensions, such as orientation, remain inactive. The neuroanatomy of vision supports this proposal, given the presence of distinct channels in the visual system that process different dimensions, such as shape, orientation, color, movement, and so forth.

When conscious images *are* constructed for a perceptual symbol, the activation of other dimensions may often be required. For example, consciously imagining a triangle may require that it have a particular orientation. [See Edelman: "Representation if Representation of Similarities" *BBS* 21(4) 1998.] However, these conscious representations need not be holistic in the sense of being irreducible to

schematic components. For example, Kosslyn and his colleagues have shown that when people construct conscious images, they construct them sequentially, component by component, not holistically in a single step (Kosslyn et al. 1988; Roth & Kosslyn 1988; see also Tye 1993).

2.2.3. Perceptual symbols need not represent specific individuals. Contrary to what some thinkers have argued, perceptual symbols need not represent specific individuals (e.g., Berkeley 1710/1982; Hume 1739/1978). Because of the schematicity assumption and its implications for human memory, we should be surprised if the cognitive system *ever* contains a complete representation of an individual. Furthermore, because of the extensive forgetting and reconstruction that characterize human memory, we should again be surprised if the cognitive system ever remembers an individual with perfect accuracy, during either conscious or unconscious processing. Typically, partial information is retrieved, and some information may be inaccurate.

As we shall see later, the *designation* of a perceptual symbol determines whether it represents a specific individual or a kind – the *resemblance* of a symbol to its referent is not critical (sect. 3.2.8). Suffice it to say for now that the same perceptual symbol can represent a variety of referents, depending on how causal and contextual factors link it to referents in different contexts (e.g., Dretske 1995; Fodor 1975; Goodman 1976; Schwartz 1981). Across different pragmatic contexts, a schematic drawing of a generic skyscraper could stand for the Empire State Building, for skyscrapers in general, or for clothing made in New York City. A drawing of the Empire State Building could likewise stand for any of these referents. Just as different physical replicas can stand for each of these referents in different contexts, perceptual representations of them can do so as well (Price 1953). Thus, the ability of a perceptual symbol to stand for a particular individual need not imply that it *must* represent an individual.

2.2.4. Perceptual symbols can be indeterminate. Theorists sometimes argue that because perceptual representations are picture-like, they are determinate. It follows that if human conceptualizations are indeterminate, perceptual representations cannot represent them (e.g., Dennett 1969; but see Block 1983). For example, it has been argued that people's conceptualizations of a tiger are indeterminate in its number of stripes; hence they must not be representing it perceptually. To my knowledge, it has not been verified empirically that people's conceptualizations of tigers *are* in fact indeterminate. If this is true, though, a perceptual representation of a tiger's stripes *could* be indeterminate in several ways (Schwartz 1981; Tye 1993). For example, the stripes could be blurred in an image, so that they are difficult to count. Or, if a perceptual symbol for stripes had been extracted schematically from the perception of a tiger, it might not contain all of the stripes but only a patch. In later representing the tiger, this free-floating symbol might be retrieved to represent the fact that the tiger was striped, but, because it was only a patch, it would not imply a particular number of stripes in the tiger. If this symbol were used to construct stripes on the surface of a simulated tiger, the tiger would then have a determinate number of stripes, but the number might differ from the original tiger, assuming for any number of reasons that the rendering of the tiger's surface did not proceed veridically.

The two solutions considered thus far assume conscious perceptual representations of a tiger. Unconscious neural representations provide another solution. It is well known that high-level neurons in perceptual systems can code information qualitatively. For example, a neuron can code the presence of a line without coding its specific length, position, or orientation. Similarly, a neuron can code the spatial frequency of a region independently of its size or location. Imagine that certain neurons in the visual system respond to stripes independently of their number (i.e., detectors for spatial frequency). In perceiving a tiger, if such detectors fire and become stored in a perceptual representation, they code a tiger's number of stripes indeterminately, because they simply respond to striped patterning and do not capture any particular number of stripes.

Qualitatively oriented neurons provide a perceptual representation system with the potential to represent a wide variety of concepts indeterminately. Consider the representation of *triangle*. Imagine that certain neurons represent the presence of lines independently of their length, position, and orientation. Further imagine that other neurons represent vertices between pairs of lines independently of the angle between them. Three qualitative detectors for lines, coupled spatially with three qualitative detectors for vertices that join them, could represent a generic triangle. Because all of these detectors are qualitative, the lengths of the lines and the angles between them do not matter; they represent all instances of *triangle* simultaneously. In this manner, qualitatively specified neurons support perceptual representations that are not only indeterminate but also generic.⁸

2.3. Multimodal perceptual symbols

The symbol formation process just described in section 2.2 can operate on any aspect of perceived experience. Not only does it operate on vision, it operates on the other four sensory modalities (audition, haptics, olfaction, and gustation), as well as on proprioception and introspection. In any modality, selective attention focuses on aspects of perceived experience and stores records of them in long-term memory, which later function as symbols. As a result, a wide variety of symbols is stored. From audition, people acquire perceptual symbols for speech and the various sounds heard in everyday experience. From touch, people acquire perceptual symbols for textures and temperatures. From proprioception, people acquire perceptual symbols for hand movements and body positions.

Presumably, each type of symbol becomes established in its respective brain area. Visual symbols become established in visual areas, auditory symbols in auditory areas, proprioceptive symbols in motor areas, and so forth. The neuroscience literature on category localization supports this assumption. When a sensory-motor area is damaged, categories that rely on it during the processing of perceived instances exhibit deficits in conceptual processing (e.g., Damasio & Damasio 1994; Gainotti et al. 1995; Pulvermüller 1999; Warrington & Shallice 1984). For example, damage to visual areas disrupts the conceptual processing of categories specified by visual features (e.g., *birds*). Analogously, damage to motor and somatosensory areas disrupts the conceptual processing of categories specified by motor and somatosensory features (e.g., *tools*). Recent neuroimaging studies on people with intact brains provide con-

verging evidence (e.g., A. Martin et al. 1995; 1996; Pulvermüller 1999; Rösler et al. 1995). When normal subjects perform conceptual tasks with *animals*, visual areas are highly active; when they perform conceptual tasks with *tools*, motor and somatosensory areas are highly active. Analogous findings have also been found for the conceptual processing of color and space (e.g., DeRenzi & Spinnler 1967; Levine et al. 1985; Rösler et al. 1995).

As these findings illustrate, perceptual symbols are multimodal, originating in all modes of perceived experience, and they are distributed widely throughout the modality-specific areas of the brain. It should now be clear that “perceptual” is not being used in its standard sense here. Rather than referring only to the sensory modalities, as it usually does, it refers much more widely to any aspect of perceived experience, including proprioception and introspection.

2.3.1. Introspection. Relative to sensory-motor processing in the brain, introspective processing is poorly understood. Functionally, three types of introspective experience appear especially important: representational states, cognitive operations, and emotional states. Representational states include the representation of an entity or event in its absence, as well as construing a perceived entity as belonging to a category. Cognitive operations include rehearsal, elaboration, search, retrieval, comparison, and transformation. Emotional states include emotions, moods, and affects. In each case, selective attention can focus on an aspect of an introspective state and stores it in memory for later use as a symbol. For example, selective attention could focus on the ability to represent something in its absence, filtering out the particular entity or event represented and storing a schematic representation of a representational state. Similarly, selective attention could focus on the process of comparison, filtering out the particular entities compared and storing a schematic representation of the comparison process. During an emotional event, selective attention could focus on emotional feelings, filtering out the specific circumstances leading to the emotion, and storing a schematic representation of the experience's “hot” components.

Much remains to be learned about the neural bases of introspection, although much is known about the neural bases of emotion (e.g., Damasio 1994; LeDoux 1996). To the extent that introspection requires attention and working memory, the neural systems that underlie them may be central (e.g., Jonides & E. Smith 1997; Posner 1995; Rushworth & Owen 1998). Like sensory-motor systems, introspection may have roots in evolution and genetics. Just as genetically constrained dimensions underlie vision (e.g., *color*, *shape*, *depth*), genetically constrained dimensions may also underlie introspection. Across individuals and cultures, these dimensions may attract selective attention, resulting in the extraction of similar perceptual symbols for introspection across individuals and cultures. Research on mental verbs in psychology and linguistics suggests what some of these dimensions might be (e.g., Cacciari & Levorato 1994; Levin 1995; Schwanenflugel et al. 1994). For example, Schwanenflugel et al. report that the dimensions of *perceptual/conceptual*, *certain/uncertain*, and *creative/noncreative* organize mental verbs such as *see*, *reason*, *know*, *guess*, and *compare*. The fact that the same dimensions arise cross-culturally suggests that different cultures conceptualize introspection similarly (e.g., D'Andrade 1987; Schwanenflugel et al., in press).

2.4. Simulators and simulations

Perceptual symbols do not exist independently of one another in long-term memory. Instead, related symbols become organized into a simulator that allows the cognitive system to construct specific simulations of an entity or event in its absence (analogous to the simulations that underlie mental imagery). Consider the process of storing perceptual symbols while viewing a particular car. As one looks at the car from the side, selective attention focuses on various aspects of its body, such as wheels, doors, and windows. As selective attention focuses on these aspects, the resulting memories are integrated spatially, perhaps using an object-centered reference frame. Similarly, as the perceiver moves to the rear of the car, to the other side, and to the front, stored perceptual records likewise become integrated into this spatially organized system. As the perceiver looks under the hood, peers into the trunk, and climbs inside the passenger area, further records become integrated. As a result of organizing perceptual records spatially, perceivers can later simulate the car in its absence. They can anticipate how the car would look from its side if they were to move around the car in the same direction as before; or they can anticipate how the car would look from the front if they were to go around the car in the opposite direction. Because they have integrated the perceptual information extracted earlier into an organized system, they can later simulate coherent experiences of the object.⁹

A similar process allows people to simulate event sequences. Imagine that someone presses the gas pedal and hears the engine roar, then lets up and hears the engine idle. Because the perceptual information stored for each subevent is not stored independently but is instead integrated temporally, the perceiver can later simulate this event sequence. Furthermore, the simulated event may contain multimodal aspects of experience, to the extent that they received selective attention. Besides visual information, the event sequence might include the proprioceptive experience of pressing the pedal, the auditory experience of hearing the engine roar, the haptic experience of feeling the car vibrating, and mild excitement about the power experienced.

As described later (sect. 2.5), the perceptual symbols extracted from an entity or event are integrated into a frame that contains perceptual symbols extracted from previous category members. For example, the perceptual symbols extracted from a car are integrated into the frame for *car*, which contains perceptual symbols extracted from previous instances. After processing many cars, a tremendous amount of multimodal information becomes established that specifies what it is like to experience cars sensorially, proprioceptively, and introspectively. In other words, the frame for *car* contains extensive multimodal information of what it is like to experience this type of thing.

A frame is never experienced directly in its entirety. Instead, subsets of frame information become active to construct specific simulations in working memory (sects. 2.4.3, 2.5.2). For example, a subset of the *car* frame might become active to simulate one particular experience of a car. On other occasions, different subsets might become active to simulate other experiences. Thus, a simulator contains two levels of structure: (1) an underlying frame that integrates perceptual symbols across category instances, and (2) the potentially infinite set of simulations that can be constructed from the frame. As we shall see in later sections,

these two levels of structure support a wide variety of important conceptual functions.¹⁰

2.4.1. Caveats. Several caveats about simulators must be noted. First, a simulator produces simulations that are *always* partial and sketchy, *never* complete. As selective attention extracts perceptual symbols from perception, it never extracts all of the information that is potentially available. As a result, a frame is impoverished relative to the perceptions that produced it, as are the simulations constructed from it.

Second, simulations are likely to be biased and distorted in various ways, rarely, if ever, being completely veridical. The well-known principles of Gestalt organization provide good examples. When a linear series of points is presented visually, an underlying line is perceived. As a result, the stored perceptual information goes beyond what is objectively present, representing a line, not just the points. The process of completion similarly goes beyond the information present. When part of an object's edge is occluded by a closer object, the edge is stored as complete, even though the entire edge was not perceived. Finally, when an imperfect edge exists on a perceived object, the perceptual system may idealize and store the edge as perfectly straight, because doing so simplifies processing. As a result, the storage of perceptual symbols may be nonveridical, as may the simulations constructed from them. McCloskey (1983) and Pylyshyn (1978) cite further distortions, and Tye (1993) suggests how perceptual simulation can explain them.

Third, a simulator is not simply an empirical collection of sense impressions but goes considerably further. Mechanisms with strong genetic constraints almost certainly play central roles in establishing, maintaining, and running simulators. For example, genetic predispositions that constrain the processing of space, objects, movement, and emotion underlie the storage of perceptual symbols and guide the simulation process (cf. Baillargeon 1995; E. Markman 1989; Spelke et al. 1992). Clearly, however, the full-blown realization of these abilities reflects considerable interaction with the environment (e.g., Elman et al. 1996). Thus, a simulator is both a "rational" and an "empirical" system, reflecting intertwined genetic and experiential histories.

2.4.2. Dispositions, schemata, and mental models. Simulators have important similarities with other constructs. In the philosophical literature, Lockean (1690/1959) dispositions and Kantian (1787/1965) schemata are comparable concepts. Both assume that unconscious generative mechanisms produce specific images of entities and events that go beyond particular entities and events experienced in the past. Similar ideas exist in more recent literatures, including Russell (1919b), Price (1953), and Damasio (1994). In all cases, two levels of structure are proposed: a deep set of generating mechanisms produces an infinite set of surface images, with the former typically being unconscious and the latter conscious.

Mental models are also related to simulators although they are not identical (K. Craik 1943; Gentner & Stevens 1983; Johnson-Laird 1983). Whereas a simulator includes two levels of structure, mental models are roughly equivalent to only the surface level, namely, simulations of specific entities and events. Mental models tend not to address underlying generative mechanisms that produce a family of related simulations.

2.4.3. Concepts, conceptualizations, and categories. According to this theory, the primary goal of human learning is to establish simulators. During childhood, the cognitive system expends much of its resources developing simulators for important types of entities and events. Once individuals can simulate a kind of thing to a culturally acceptable degree, they have an adequate understanding of it. What is deemed a culturally competent grasp of a category may vary, but in general it can be viewed as the ability to simulate the range of multimodal experiences common to the majority of a culture's members (cf. Romney et al. 1986). Thus, people have a culturally acceptable simulator for *chair* if they can construct multimodal simulations of the chairs typically encountered in their culture, as well as the activities associated with them.

In this theory, a *concept* is equivalent to a simulator. It is the knowledge and accompanying processes that allow an individual to represent some kind of entity or event adequately. A given simulator can produce limitless simulations of a kind, with each simulation providing a different *conceptualization* of it. Whereas a concept represents a kind generally, a conceptualization provides one specific way of thinking about it. For example, the simulator for *chair* can simulate many different chairs under many different circumstances, each comprising a different conceptualization of the category. For further discussion of this distinction between permanent knowledge of a kind in long-term memory and temporary representations of it in working memory, see Barsalou (1987; 1989; 1993; also see sect. 2.4.5).

Simulators do not arise in a vacuum but develop to track meaningful units in the world. As a result, knowledge can accumulate for each unit over time and support optimal interactions with it (e.g., Barsalou et al. 1993; 1998; Millikan 1998). Meaningful units include important individuals (e.g., family members, friends, personal possessions) and categories (e.g., natural kinds, artifacts, events), where a *category* is a set of individuals in the environment or introspection. Once a simulator becomes established in memory for a category, it helps identify members of the category on subsequent occasions, and it provides categorical inferences about them, as described next.¹¹

2.4.4. Categorization, categorical inferences, and affordances. Tracking a category successfully requires that its members be categorized correctly when they appear. Viewing concepts as simulators suggests a different way of thinking about *categorization*. Whereas many theories assume that relatively static, amodal structures determine category membership (e.g., definitions, prototypes, exemplars, theories), simulators suggest a more dynamic, embodied approach: if the simulator for a category can produce a satisfactory simulation of a perceived entity, the entity belongs in the category. If the simulator cannot produce a satisfactory simulation, the entity is not a category member.¹²

Besides being dynamic, grounding categorization in perceptual symbols has another important feature: the knowledge that determines categorization is represented in roughly the same manner as the perceived entities that must be categorized. For example, the perceptual simulations used to categorize chairs approximate the actual perceptions of chairs. In contrast, amodal theories assume that amodal features in concepts are compared to perceived entities to perform categorization. Whereas amodal theories have to explain how two very different types of representa-

tion are compared, perceptual symbol systems simply assume that two similar representations are compared. As a natural side effect of perceiving a category's members, perceptual knowledge accrues that can be compared directly to perceived entities during categorization.

On this view, categorization depends on both familiar and novel simulations. Each successful categorization stores a simulation of the entity categorized. If the same entity or a highly similar one is encountered later, it is assigned to the category because the perception of it matches an existing simulation in memory. Alternatively, if a novel entity is encountered that fails to match an existing simulation, constructing a novel simulation that matches the entity can establish membership. Explanation-based learning assumes a similar distinction between expertise and creativity in categorization (DeJong & Mooney 1986; T. Mitchell et al. 1986), as do theories of skill acquisition (Anderson 1993; Logan 1988; Newell 1990), although these approaches typically adopt amodal representations.

As an example, imagine that the simulator for *triangle* constructs three lines and connects their ends uniquely. Following experiences with previous triangles, simulations that match these instances become stored in the simulator. On encountering these triangles later, or highly similar ones, prestored simulations support rapid categorization, thereby implementing expertise. However, a very different triangle, never seen before, can also be categorized if the *triangle* simulator can construct a simulation of it (cf. Miller & Johnson-Laird 1976).

Categorization is not an end in itself but provides access to *categorical inferences*. Once an entity is categorized, knowledge associated with the category provides predictions about the entity's structure, history, and behavior, and also suggests ways of interacting with it (e.g., Barsalou 1991; Ross 1996; see also sect. 3.2.2). In this theory, categorical inferences arise through simulation. Because a simulator contains a tremendous amount of multimodal knowledge about a category, it can simulate information that goes beyond that perceived in a categorized entity. On perceiving a computer from the front, the simulator for *computer* can simulate all sorts of things not perceived, such as the computer's rear panel and internal components, what the computer will do when turned on, what tasks it can perform, how the keys will feel when pressed, and so forth. Rather than having to learn about the entity from scratch, a perceiver can run simulations that anticipate the entity's structure and behavior and that suggest ways of interacting with it successfully.

Simulators also produce categorical inferences in the absence of category members. As described later, simulations provide people with a powerful ability to reason about entities and events in their absence (sects. 2.6, 3.1, 3.2, 4.1, 4.2). Simulations of future entities, such as a rental home, allow people to identify preparations that must be made in advance. Simulations of future events, such as asking a favor, allow people to identify optimal strategies for achieving success. To the extent that future category members are similar to previous category members, simulations of previous members provide reasonable inferences about future members.¹³

Deriving categorical inferences successfully requires that simulations preserve at least some of the *affordances* present in actual sensory-motor experiences with category members (cf. Gibson 1979; also see S. Edelman 1998). To

the extent that simulations capture affordances from perception and action, successful reasoning about physical situations can proceed in their absence (Glenberg 1997; Glenberg et al. 1998b; Newton 1996). Agents can draw inferences that go beyond perceived entities, and they can plan intelligently for the future. While sitting in a restaurant and wanting to hide from someone entering, one could simulate that a newspaper on the table affords covering one's face completely but that a matchbook does not. As a result of these simulations, the newspaper is selected to achieve this goal rather than the matchbook. Because the simulations captured the physical affordances correctly, the selected strategy works.

2.4.5. Concept stability. Equating concepts with simulators provides a solution to the problem of concept stability. Previous work demonstrates that conceptualizations of a category vary widely between and within individuals (Barsalou 1987; 1989; 1993). If different people conceptualize *bird* differently on a given occasion, and if the same individual conceptualizes *bird* differently across occasions, how can stability be achieved for this concept?

One solution is to assume that a common simulator for *bird* underlies these different conceptualizations, both between and within individuals. First, consider how a simulator produces stability within an individual. If a person's different simulations of a category arise from the same simulator, then they can all be viewed as instantiating the same concept. Because the same simulator produced all of these simulations, it unifies them. Between individuals, the key issue concerns whether different people acquire similar simulators. A number of factors suggest that they should, including a common cognitive system, common experience with the physical world, and socio-cultural institutions that induce conventions (e.g., Newton 1996; Tomasello et al. 1993). Although two individuals may represent the same category differently on a given occasion, each may have the ability to simulate the other's conceptualization. In an unpublished study, subjects almost always viewed other subjects' conceptualizations of a category as correct, even though their individual conceptualizations varied widely. Each subject produced a unique conceptualization but accepted those of other subjects because they could be simulated. Furthermore, common contextual constraints during communication often drive two people's simulations of a category into similar forms. In another unpublished study, conceptualizations of a category became much more stable both between and within subjects when constructed in a common context. Subjects shared similar simulators that produced similar conceptualizations when constrained adequately.

2.4.6. Cognitive penetration. The notion of a simulator is difficult to reconcile with the view that cognition does not penetrate perception (Fodor 1983). According to the impenetrability hypothesis, the amodal symbol system underlying higher cognition has little or no impact on processing in sensory-motor systems because these systems are modular and therefore impenetrable. In contrast, the construct of a simulator assumes that sensory-motor systems are deeply penetrable. Because perceptual symbols reside in sensory-motor systems, running a simulator involves a partial running of these systems in a top-down manner.

In an insightful *BBS* review of top-down effects in vision,

Pylyshyn (1999) concludes that cognition only produces top-down effects indirectly through attention and decision making – it does not affect the content of vision directly. Contrary to this conclusion, however, much evidence indicates that cognition *does* affect the content of sensory-motor systems directly. The neuroscience literature on mental imagery demonstrates clearly that cognition establishes content in sensory-motor systems in the absence of physical input. In visual imagery, the primary visual cortex, V1, is often active, along with many other early visual areas (e.g., Kosslyn et al. 1995). In motor imagery, the primary motor cortex, M1, is often active, along with many other early motor areas (e.g., Crammond 1997; Deschaumes-Molinario et al. 1992; Jeannerod 1994; 1995). Indeed, motor imagery not only activates early motor areas, it also stimulates spinal neurons, produces limb movements, and modulates both respiration and heart rate. When sharpshooters imagine shooting a gun, their entire body behaves similarly to actually doing so. In auditory imagery, activation has not yet been observed in the primary auditory cortex, but activation has been observed in other early auditory areas (e.g., Zatorre et al. 1996). These findings clearly demonstrate that cognition establishes content in sensory-motor systems in the absence of physical input.

A potential response is that mental imagery arises solely within sensory-motor areas – it is not initiated by cognitive areas. In this vein, Pylyshyn (1999) suggests that perceptual modules contain local memory areas that affect the content of perception in a top-down manner. This move undermines the impenetrability thesis, however, at least in its strong form. As a quick perusal of textbooks on cognition and perception reveals, memory is widely viewed as a basic cognitive process – not as a perceptual process. Many researchers would probably agree that once memory is imported into a sensory-motor system, cognition has been imported. Furthermore, to distinguish perceptual memory from cognitive memory, as Pylyshyn does, makes sense only if one assumes that cognition utilizes amodal symbols. Once one adopts the perspective of perceptual symbol systems, there is only perceptual memory, and it constitutes the representational core of cognition. In this spirit, Damasio (1989) argues eloquently that there is no sharp discontinuity between perceptual and cognitive memory. Instead, there is simply a gradient from posterior to anterior association areas in the complexity and specificity of the memories that they activate in sensory-motor areas. On Damasio's view, memory areas both inside and outside a sensory-motor system control its feature map to implement cognitive representations. In this spirit, the remainder of this target article assumes that top-down cognitive processing includes all memory effects on perceptual content, including memory effects that originate in local association areas.

Nevertheless, Pylyshyn (1999) makes compelling arguments about the resiliency of bottom-up information in face-to-face competition with contradicting top-down information. For example, when staring at the Müller-Lyer illusion, one cannot perceive the horizontal lines as equivalent in length, even though one knows cognitively that they are. Rather than indicating impenetrability, however, this important observation may simply indicate that bottom-up information dominates top-down information when they conflict (except in the degenerate case of psychosis and other hallucinatory states, when top-down information *does* dominate bottom-up information). Indeed, Marslen-Wil-

son and Tyler (1980), although taking a nonmodular interactive approach, offer exactly this account of bottom-up dominance in speech recognition. Although top-down processing can penetrate speech processing, it is overridden when bottom-up information conflicts. If semantic and syntactic knowledge predict that “The cowboy climbed into the _____” ends with “saddle,” but the final word is actually “jacuzzi,” then “jacuzzi” overrides “saddle.”

On this view, sensory-motor systems are penetrable but not always. When bottom-up information conflicts with top-down information, the former usually dominates. When bottom-up information is absent, however, top-down information penetrates, as in mental imagery. Perhaps most critically, when bottom-up and top-down information are compatible, top-down processing again penetrates, but in subtle manners that complement bottom-up processing. The next section (sect. 2.4.7) reviews several important phenomena in which bottom-up and top-down processing simultaneously activate sensory-motor representations as they cooperate to perceive physical entities (i.e., implicit memory, filling-in, anticipation, interpretation). Recent work on simultaneous imagery and perception shows clearly that these two processes work well together when compatible (e.g., Craver-Lemley & Reeves 1997; Gilden et al. 1995).

Perhaps the critical issue in this debate concerns the definition of cognition. On Pylyshyn’s view, cognition concerns semantic beliefs about the external world (i.e., the belief that the horizontal lines are the same length in the Müller-Lyer illusion). However, this is a far narrower view of cognition than most cognitive psychologists take, as evidenced by journals and texts in the field. Judging from these sources, cognitive psychologists believe that a much broader array of processes and representations – including memory – constitutes cognition.

Ultimately, as Pylyshyn suggests, identifying the mechanisms that underlie intelligence should be our primary goal, from the most preliminary sensory processes to the most abstract thought processes. Where we actually draw the line between perception and cognition may not be all that important, useful, or meaningful. In this spirit, perceptual symbol systems attempt to characterize the mechanisms that underlie the human conceptual system. As we have seen, the primary thesis is that sensory-motor systems represent not only perceived entities but also conceptualizations of them in their absence. From this perspective, cognition penetrates perception when sensory input is absent, or when top-down inferences are compatible with sensory input.

2.4.7. A family of representational processes. Evolution often capitalizes on existing mechanisms to perform new functions (Gould 1991). Representational mechanisms in sensory-motor regions of the brain may be such mechanisms. Thus far, these representational mechanisms have played three roles in perceptual symbol systems: (1) In perception, they represent physical objects. (2) In imagery, they represent objects in their absence. (3) In conception, they also represent objects in their absence. On this view, conception differs from imagery primarily in the consciousness and specificity of sensory-motor representations, with these representations being more conscious and detailed in imagery than in conception (Solomon & Barsalou 1999a; Wu & Barsalou 1999). Several other cognitive

processes also appear to use the same representational mechanisms, including implicit memory, filling-in, anticipation, and interpretation. Whereas perception, imagery, and conception perform either bottom-up or top-down processing exclusively, these other four processes fuse complementary *mixtures* of bottom-up and top-down processing to construct perceptions.

In implicit memory (i.e., repetition priming), a perceptual memory speeds the perception of a familiar entity (e.g., Roediger & McDermott 1993; Schacter 1995). On seeing a particular chair, for example, a memory is established that speeds perception of the same chair later. Much research demonstrates the strong perceptual character of these memories, with slight deviations in perceptual features eliminating facilitative effects. Furthermore, imagining an entity produces much the same facilitation as perceiving it, suggesting a common representational basis. Perhaps most critically, implicit memory has been localized in sensory-motor areas of the brain, with *decreasing* brain activity required to perceive a familiar entity (e.g., Buckner et al. 1995). Thus, the representations that underlie implicit memory reside in the same systems that process entities perceptually. When a familiar entity is perceived, implicit memories become fused with bottom-up information to represent it efficiently.

In filling-in, a perceptual memory completes gaps in bottom-up information. Some filling-in phenomena reflect perceptual inferences that are largely independent of memory (for a review, see Pessoa et al. 1998). In the perception of illusory contours, for example, low-level sensory mechanisms infer edges on the basis of perceived vertices (e.g., Kanizsa 1979). However, other filling-in phenomena rely heavily on memory. In the phoneme restoration effect, knowledge of a word creates the conscious perceptual experience of hearing a phoneme where noise exists physically (e.g., Samuel 1981; 1987; Warren 1970). More significantly, phoneme restoration adapts low-level feature detectors much as if physical phonemes had adapted them (Samuel 1997). Thus, when a word is recognized, its memory representation fills in missing phonemes, not only in experience, but also in sensory processing. Such findings strongly suggest that cognitive and perceptual representations reside in a common system, and that they become fused to produce perceptual representations. Knowledge-based filling-in also occurs in vision. For example, knowledge about bodily movements causes apparent motion to deviate from the perceptual principle of minimal distance (Shiffrar & Freyd 1990; 1993). Rather than filling in an arm as taking the shortest path through a torso, perceivers fill it in as taking the longer path around the torso, consistent with bodily experience.

In perceptual anticipation, the cognitive system uses past experience to simulate a perceived entity’s future activity. For example, if an object traveling along a trajectory disappears, perceivers anticipate where it would be if it were still on the trajectory, recognizing it faster at this point than at the point it disappeared, or at any other point in the display (Freyd 1987). Recent findings indicate that knowledge affects the simulation of these trajectories. When subjects believe that an ambiguous object is a rocket, they simulate a different trajectory compared to when they believe it is a steeple (Reed & Vinson 1996). Even infants produce perceptual anticipations in various occlusion tasks (e.g., Bailargeon 1995; Hespos & Rochat 1997). These results fur-

ther indicate that top-down and bottom-up processes coordinate the construction of useful perceptions.

In interpretation, the conceptual representation of an ambiguous perceptual stimulus biases sensory processing. In audition, when subjects believe that multiple speakers are producing a series of speech sounds, they normalize the sounds differently for each speaker (Magnuson & Nusbaum 1993). In contrast, when subjects believe that only one speaker is producing the same sounds, they treat them instead as differences in the speaker's emphasis. Thus, each interpretation produces sensory processing that is appropriate for its particular conceptualization of the world. Again, cognition and sensation coordinate to produce meaningful perceptions (see also Nusbaum & Morin 1992; Schwab 1981). Analogous interpretive effects occur in vision. Conceptual interpretations guide computations of figure and ground in early visual processing (Peterson & Gibson 1993; Weisstein & Wong 1986); they affect the selective adaptation of spatial frequency detectors (Weisstein et al. 1972; Weisstein & Harris 1980); and they facilitate edge detection (Weisstein & Harris 1974). Frith and Dolan (1997) report that top-down interpretive processing activates sensory-motor regions in the brain.

In summary, an important family of basic cognitive processes appears to utilize a single mechanism, namely, sensory-motor representations. These processes, although related, vary along a continuum of bottom-up to top-down processing. At one extreme, bottom-up input activates sensory-motor representations in the absence of top-down processing ("pure" perception). At the other extreme, top-down processing activates sensory-motor representations in the absence of bottom-up processing (imagery and conception). In between lie processes that fuse complementary mixtures of bottom-up and top-down processing to coordinate the perception of physical entities (implicit memory, filling-in, anticipation, interpretation).

2.5. Frames

A frame is an integrated system of perceptual symbols that is used to construct specific simulations of a category. Together, a frame and the simulations it produces constitute a simulator. In most theories, frames are amodal, as are the closely related constructs of schemata and scripts (e.g., Minsky 1977; 1985; Rumelhart & Ortony 1978; Schank & Abelson 1977; for reviews, see Barsalou 1992; Barsalou & Hale 1992). As we shall see, however, frames and schemata have natural analogues in perceptual symbol systems. In the account that follows, all aspects require much further development, and many important issues are not addressed. This account is meant to provide only a rough sense of how frames develop in perceptual symbol systems and how they produce specific simulations.

The partial frame for *car* in Figure 3 illustrates how perceptual symbol systems implement frames. On the perception of a first car (Fig. 3A), the schematic symbol formation process in section 2.2 extracts perceptual symbols for the car's overall shape and some of its components, and then integrates these symbols into an object-centered reference frame. The representation at the top of Figure 3A represents the approximate volumetric extent of the car, as well as embedded subregions that contain significant components (e.g., the doors and tires). These subregions and their specializations reflect the result of the symbol formation




Figure 3. (A) An example of establishing an initial frame for *car* after processing a first instance. (B) The frame's evolution after processing a second instance. (C) Constructing a simulation of the second instance from the frame in panel B. In Panels A and B, lines with arrows represent excitatory connections, and lines with circles represent inhibitory connections.

process described earlier. For every subregion that receives selective attention, an approximate delineation of the subregion is stored and then connected to the content information that specializes it. For example, attending to the window and handle of the front door establishes perceptual symbols for these subregions and the content information that specializes them.¹⁴

As Figure 3A illustrates, the frame represents spatial and content information separately. At one level, the volumetric regions of the object are represented according to their spatial layout. At another level, the contents of these subregions are represented as specializations. This distinction between levels of representation follows the work of Ungerleider and Mishkin (1982), who identified separate neural pathways for spatial/motor information and object features (also see Milner & Goodale 1995). Whereas the spatial representation establishes the frame's skeleton, the content specializations flesh it out.

On the perception of a second car, a reminding takes place, retrieving the spatially integrated set of symbols for the first car (Barsalou et al. 1998; Medin & Ross 1989; Millikan 1998; Ross et al. 1990; Spalding & Ross 1994). The retrieved set of symbols guides processing of the second car in a top-down manner, leading to the extraction of perceptual symbols in the same subregions. As Figure 3B illustrates, this might lead to the extraction of content information for the second car's shape, doors, and wheels, which become connected to the same subregions as the content extracted from the first car. In addition, other subregions of the second car may be processed, establishing new perceptual symbols (e.g., the antenna and gas cap).

Most important, all the information extracted from the two cars becomes integrated into a knowledge structure that constitutes the beginnings of a *car* frame. When perceptual symbols have been extracted from the same volumetric region for both cars, they both become connected to that subregion in the object-oriented reference frame. For example, the perceptual symbols for each car's doors become associated with the *door* subregions. As a result, the specialization from either car can be retrieved to specialize a region during a simulation, as may an average superimposition of them.

Figure 3B further illustrates how connections within the frame change as new instances become encoded. First, specializations from the same instance are connected together, providing a mechanism for later reinstating it. Second, these connections become weaker over time, with the dashed connections for the first instance representing weaker connections than those for the second instance. Third, inhibitory connections develop between specializations that compete for the same subregion, providing an additional mechanism for reinstating particular instances. Finally, connections processed repeatedly become stronger, such as the thicker connections between the overall volume and the subregions for the doors and wheels.

Following experiences with many cars, the *car* frame accrues a tremendous amount of information. For any given subregion, many specializations exist, and default specializations develop, either through the averaging of specializations, or because one specialization develops the strongest association to the subregion. Frames develop similarly for event concepts (e.g., *eating*), except that subregions exist over time as well as in space (examples will be presented in sect. 3.4 on abstract concepts).

2.5.1. Basic frame properties. Barsalou (1992) and Barsalou and Hale (1993) propose that four basic properties constitute a frame (and the related construct of a schema): (1) predicates, (2) attribute-value bindings, (3) constraints, and (4) recursion. Barsalou (1993) describes how these properties manifest themselves in perceptual symbol systems.

Predicates are roughly equivalent to unspecialized frames. For example, the predicate $CAR(Door = x, Window = y, \dots)$ is roughly equivalent to the perceptual frame for *car* with its subregions unspecialized (e.g., the hierarchical volumetric representation in Fig. 3B).

Attribute-value bindings arise through the specialization of a subregion in a simulation. As different specializations become bound to the same subregion, they establish different "values" of an "attribute" or "slot" (e.g., *door* and its specializations in Fig. 3B).

Constraints arise from associative connections between specializations that represent individuals and subcategories within a frame (e.g., the connections between the specializations for the second instance of *car* in Fig. 3B). Thus, activating the second car's specialization in one subregion activates its associated specializations in other subregions, thereby simulating this specific car. Because constraints can weaken over time and because strong defaults may dominate the activation process, reconstructive error can occur. For example, if the first car's tires received extensive processing, the strength of their connections to the subregions for the wheels might dominate, even during attempts to simulate the second car.

Recursion arises from establishing a simulation within an

existing simulator. As Figures 3A and 3B suggest, the simulator for *wheel* might initially just simulate schematic circular volumes, with all other information filtered out. As more attention is paid to the wheels on later occasions, however, detail is extracted from their subregions. As a result, simulators for *tire* and *hubcap* develop within the simulator for *wheel*, organizing the various specializations that occur for each of them. Because this process can continue indefinitely, an arbitrarily deep system of simulators develops, bounded only by the perceiver's motivation and perceptual resolution.

2.5.2. Constructing specific simulations. The cognitive system uses the frame for a category to construct specific simulations (i.e., mental models). As discussed in later sections, such simulations take myriad forms. In general, however, the process takes the form illustrated in Figure 3C. First, the overall volumetric representation for *car* becomes active, along with a subset of its subregions (e.g., for doors, wheels, antenna, gas cap). If a cursory simulation is being constructed, only the subregions most frequently processed previously are included, with other subregions remaining inactive. This process can then proceed recursively, with further subregions and specializations competing for inclusion. As Figure 3C illustrates, subregions for the doors and windows subsequently become active, followed by their specializations.

In any given subregion, the specialization having the highest association with the subregion, with other active regions, and with other active specializations becomes active. Because this process occurs for each subregion simultaneously, it is highly interactive. The simulation that emerges reflects the strongest attractor in the frame's state space. If context "clamps" certain aspects of the simulation initially, the constraint satisfaction process may diverge from the frame's strongest attractor toward a weaker one. If an event is being simulated, subregions and their specializations may change over time as the constraint satisfaction process evolves recurrently.

During a simulation, processing is not limited to the retrieval of frame information but can also include transformations of it. Retrieved information can be enlarged, shrunk, stretched, and reshaped; it can be translated across the simulation spatially or temporally; it can be rotated in any dimension; it can remain fixed while the perspective on it varies; it can be broken into pieces; it can be merged with other information. Other transformations are no doubt possible as well. The imagery literature offers compelling evidence that such transformations are readily available in the cognitive system (e.g., Finke 1989; Kosslyn 1980; Shepard & Cooper 1982), and that these transformations conform closely to perceptual experience (e.g., Freyd 1987; Parsons 1987a; 1987b; Shiffrar & Freyd 1990; 1993).

2.5.3. Framing and background-dependent meaning. As linguists and philosophers have noted for some time, concepts are often specified relative to one another (e.g., Fillmore 1985; Langacker 1986; A. Lehrer 1974; A. Lehrer & Kittay 1992; Quine 1953). Psychologists often make a similar observation that concepts are specified in the context of intuitive theories (e.g., Carey 1985; Keil 1989; Murphy & Medin 1985; Rips 1989; Wellman & Gelman 1988). Framing and background-dependent meaning are two specific ways in which background knowledge specifies concepts. In

framing, a focal concept depends on a background concept and cannot be specified independently of it. For example, *payment* is specified relative to *buy*, *hypotenuse* is specified relative to *right triangle*, and *foot* is specified relative to *leg* and *body*. In background-dependent meaning, the same focal concept changes as its background concept changes. For example, the conceptualization of *foot* varies as the background concept changes from *human* to *horse* to *tree*. Similarly, the conceptualization of *handle* varies across *shovel*, *drawer*, and *car door*; as does the conceptualization of *red* across *fire truck*, *brick*, *hair*, and *wine* (Half et al. 1976).

Frames provide the background structures that support framing. Thus, the event frame for *buy* organizes the background knowledge necessary for understanding *payment*, and the entity frame for *human* organizes the background knowledge necessary for understanding *foot*. When *payment* or *foot* is conceptualized, its associated frame produces a simulation that provides the necessary background for understanding it. For example, a simulation of a human body provides one possible framing for a simulation of a foot.

Frames also offer a natural account of background dependent meaning. *Foot*, for example, is conceptualized differently when *human* is simulated in the background than when *horse* or *tree* is simulated. Because different perceptual symbols are accessed for *foot* in the context of different frames, simulations of *foot* vary widely. Similarly, different conceptualizations of *red* reflect different perceptual symbols accessed in frames for *fire truck*, *brick*, *hair*, and *wine*.¹⁵

2.6. Linguistic indexing and control

In humans, linguistic symbols develop together with their associated perceptual symbols. Like a perceptual symbol, a linguistic symbol is a schematic memory of a perceived event, where the perceived event is a spoken or a written word. A linguistic symbol is *not* an amodal symbol, nor does an amodal symbol *ever* develop in conjunction with it. Instead, a linguistic symbol develops just like a perceptual symbol. As selective attention focuses on spoken and written words, schematic memories extracted from perceptual states become integrated into simulators that later produce simulations of these words in recognition, imagination, and production.

As simulators for words develop in memory, they become associated with simulators for the entities and events to which they refer. Whereas some simulators for words become linked to simulators for entire entities or events, others become linked to subregions and specializations. Whereas “car” becomes linked to the entire simulator for *car*, “trunk” becomes linked to one of its subregions. Simulators for words also become associated with other aspects of simulations, including surface properties (e.g., “red”), manners (e.g., “rapidly”), relations (e.g., “above”), and so forth. Within the simulator for a concept, large numbers of simulators for words become associated with its various aspects to produce a semantic field that mirrors the underlying conceptual field (Barsalou 1991; 1992; 1993).

Once simulators for words become linked to simulators for concepts, they can control simulations. On recognizing a word, the cognitive system activates the simulator for the associated concept to simulate a possible referent. On parsing the sentences in a text, surface syntax provides instruc-

tions for building perceptual simulations (Langacker 1986; 1987; 1991; 1997). As discussed in the next section, the productive nature of language, coupled with the links between linguistic and perceptual simulators, provides a powerful means of constructing simulations that go far beyond an individual’s experience. As people hear or read a text, they use productively formulated sentences to construct a productively formulated simulation that constitutes a semantic interpretation (sect. 4.1.6). Conversely, during language production, the construction of a simulation activates associated words and syntactic patterns, which become candidates for spoken sentences designed to produce a similar simulation in a listener. Thus, linguistic symbols index and control simulations to provide humans with a conceptual ability that is probably the most powerful of any species (Donald 1991; 1993). As MacWhinney (1998) suggests, language allows conversationalists to coordinate simulations from a wide variety of useful perspectives.¹⁶

3. Derived properties

By parsing perception into schematic components and then integrating components across individuals into frames, simulators develop that represent the types of entities and events in experience. The result is a basic conceptual system, not a recording system. Rather than recording holistic representations of perceptual experience, this system establishes knowledge about the categories of individuals that constitute it. Each simulator represents a type, not a token, and the collection of simulators constitutes a basic conceptual system that represents the components of perception and provides categorical inferences about them.

We have yet to characterize a fully functional conceptual system. To do so, we must establish that perceptual symbol systems can implement productivity, propositions, and abstract concepts. The next section demonstrates that these properties follow naturally from the basic system presented thus far and that one additional property – variable embodiment – follows as well.

3.1. Productivity

One of the important lessons we have learned from amodal symbol systems is that a viable theory of knowledge must be productive (e.g., Chomsky 1957; Fodor 1975; Fodor & Pylyshyn 1988). The human cognitive system can produce an infinite number of conceptual and linguistic structures that go far beyond those experienced. No one has experienced a real Cheshire Cat, but it is easy to imagine a cat whose body fades and reappears while its human smile remains.

Productivity is the ability to construct an unlimited number of complex representations from a finite number of symbols using combinatorial and recursive mechanisms. It is fair to say that productivity is not typically recognized as possible within perceptual theories of knowledge, again because these theories are usually construed as recording systems. It is worth noting, however, that certain physical artifacts exhibit a kind of “perceptual productivity,” using combinatorial and recursive procedures on schematic diagrams to construct limitless complex diagrams. In architecture, notations exist for combining primitive schematic diagrams combinatorially and recursively to form complex diagrams of buildings (e.g., W. Mitchell 1990). In electron-

ics, similar notations exist for building complex circuits from primitive schematic components (Haugeland 1991). The fact that physical artifacts can function in this manner suggests that schematic perceptual representations in cognition could behave similarly (Price 1953). As we shall see, if one adopts the core properties of perceptual symbol systems established thus far, productivity follows naturally.

Figure 4 illustrates productivity in perceptual symbol systems. Before exploring productivity in detail, it is first necessary to make several points about the diagrams in Figure 4, as well as those in all later figures. First, diagrams such as the balloon in Figure 4A should *not* be viewed as literally representing pictures or conscious images. Instead, these theoretical illustrations stand for configurations of neurons that become active in representing the physical information conveyed in these drawings. For example, the balloon in Figure 4A is a theoretical notation that refers to configurations of neurons active in perceiving balloons.

Second, each of these drawings stands metonymically for a simulator. Rather than standing for only one projection of an object (e.g., the balloon in Fig. 4A), these drawings stand for a simulator capable of producing countless projections of the instance shown, as well as of an unlimited number of other instances.

Third, the diagrams in Figure 4B stand for simulators of spatial relations that result from the symbol formation process described earlier. During the perception of a balloon above a cloud, for example, selective attention focuses on the occupied regions of space, filtering out the entities in them. As a result, a schematic representation of *above* develops that contains two schematic regions of space within one of several possible reference frames (not shown in Fig. 4). Following the similar extraction of information on other occasions, a simulator develops that can render many different *above* relations. For example, specific simulations may vary in the vertical distance between the two regions, in their horizontal offset, and so forth. Finally, the thicker boundary for a given region in a spatial relation indicates that selective attention is focusing on it.¹⁷ Thus, *above* and *below* involve the same representation of space, but with a different distribution of attention over it. As much research illustrates, an adequate treatment of spatial concepts requires further analysis and detail than provided here (e.g., Herskovits 1997; Regier 1996; Talmy 1983). Nevertheless, one can view spatial concepts as simulators that develop through the schematic symbol formation process in section 2.2.

Figure 4C illustrates how the simulators in Figures 4A and 4B combine to produce complex perceptual simulations combinatorially. In the leftmost example, the simulator for *above* produces a specific simulation of *above* (i.e., two schematic regions of the same size, close together and in vertical alignment). The simulators for *balloon* and *cloud* produce specific simulations that specialize the two regions of the *above* simulation. The result is a complex simulation in which a balloon is simulated above a cloud. Note that complex simulations of these three categories could take infinitely many forms, including different balloons, clouds, and above relations. The second and third examples in Figure 4C illustrate the combinatorial nature of these simulations, as the various objects in Figure 4A are rotated through each of the regions in *above*. Because many possible objects could enter into simulations of *above*, a very large number of such simulations is possible.




Figure 4. An example of how perceptual symbols for object categories (A) and spatial relations (B) implement productivity through combinatorial (C) and recursive (D) processing. Boxes with thin solid lines represent simulators; boxes with thick dashed lines represent simulations.

A simulation can be constructed recursively by specializing the specialization of a schematic region. In Figure 4D, the lower region of *above* is specialized recursively with a simulation of *left-of*, whose regions are then specialized with simulations of *jet* and *cloud*. The resulting simulation represents a balloon that is above a jet to the left of a cloud. Because such recursion can occur indefinitely, an infinite number of simulations can be produced in principle.

3.1.1. Reversing the symbol formation process. Productivity in perceptual symbol systems is approximately the symbol formation process run in reverse. During symbol formation, large amounts of information are filtered out of perceptual representations to form a schematic representation of a selected aspect (sect. 2.2). During productivity, small amounts of the information filtered out are added back. Thus, schematicity makes productivity possible. For example, if a perceptual symbol for *ball* only represents its shape schematically after color and texture have been filtered out, then information about color and texture can later be added productively. For example, the simulation of a *ball* could evolve into a *blue ball* or a *smooth yellow ball*. Because the symbol formation process similarly establishes schematic representations for colors and textures, these representations can be combined productively with perceptual representations for shapes to produce complex simulations.

Productivity is not limited to filling in schematic regions

but can also result from replacements, transformations, and deletions of existing structure. Imagine that a simulated lamp includes a white cylindrical shade. To represent different lamps, the simulated shade could be replaced with a simulated cardboard box, it could be transformed to a cone shape, or it could be deleted altogether. Such operations appear widely available for constructing simulations, extending productivity further.

Most important, the complementarity of schematization and specialization allow perceptual systems to go beyond a recording system and become a conceptual system. Whereas photos and videos only capture information holistically, a perceptual symbol system extracts particular parts of images schematically and integrates them into simulators. Once simulators exist, they can be combined to construct simulations productively. Such abilities go far beyond the recording abilities of photos and videos; yet, as we have seen, they can be achieved within a perceptual framework.

3.1.2. Productivity in imagination. Productivity can surpass experience in many ways, constituting an important source of creativity (Barsalou & Prinz 1997). For example, one can simulate a chair never encountered, such as a *pitted lavender chair*. Because perceptual symbols for colors become organized together in a semantic field, as do perceptual symbols for textures, simulations of a chair can cycle through the symbols within each field to try out various combinations (Barsalou 1991; 1992; 1993). During interior decoration, one can combine colors and textures productively with a schematic representation of a chair to see which works best, with the semantic fields for colors and textures providing “palettes” for constructing the possibilities. Because different semantic fields can be explored orthogonally, the resulting simulations take on an analytic combinatorial character.

Productive mechanisms can further construct simulations that violate properties of the physical world. For example, one can productively combine the shape of a chair with the simulation of a dog to construct a dog that functions as a chair. By searching through the combinatorial space of possibilities, one can construct many similar simulations, such as Carroll’s (1960) flamingos in *Alice in Wonderland* who form themselves into croquet mallets, and his hedgehogs who form themselves into croquet balls. The space of such possibilities is very large, bounded only by the ranges of animals and artifacts that can be combined productively. When the recursive possibilities are considered, the possibilities become infinite (e.g., a camel taking the form of a carriage, with alligators that form themselves into wheels, with starfish that form themselves into spokes, etc.). Children’s books are full of many further examples, such as productively combining various human abilities with various kinds of animals (e.g., dinosaurs that talk, tall birds that build sand castles).

As these examples illustrate, the human conceptual ability transcends experience, combining existing knowledge in new ways. Because perceptual symbols are schematic, they can combine creatively to simulate imaginary entities. Wu and Barsalou (1999) demonstrate that when people form novel concepts productively, perceptual simulation plays a central role.

3.1.3. Constraints and emergent properties. Although the productive potential of perceptual symbols is extensive, it is

not a simple process whereby any two perceptual symbols can combine to form a whole that is equal to the sum of its parts. Instead, there are constraints on this process, as well as emergent features (Prinz 1997; Prinz & Barsalou, in press a). Presumably, these constraints and emergent features reflect affordances captured through the schematic symbol formation process (sect. 2.4.4).

Constraints arise when a schematic perceptual symbol cannot be applied to a simulated entity, because the simulation lacks a critical characteristic. For example, it is difficult to transform a simulated *watermelon* into a *running watermelon*, because an entity requires legs in order to run. If a simulated entity does not have legs, then simulating it running is difficult. Interestingly, even if an entity has the wrong kind of legs, such as a *chair*; it can easily be simulated as running, because it has the requisite spatial parts that enable the transformation.

Emergent properties arise frequently during the productive construction of simulations. As Langacker (1987) observes, combining animate agents productively with *running* produces emergent properties with the methods and manners of running varying across humans, birds, horses, crabs, and spiders. Although emergent properties may often reflect perceptual memories of familiar animals running (cf. Freyd 1987; Parsons 1987a; 1987b; Reed & Vinson 1996; Shiffrar & Freyd 1990; 1993), they may also arise from spatiotemporal properties of the simulation process. For example, when one imagines a chair versus a sofa running, different simulations result. Because people have probably never seen a chair or a sofa run, these different simulations probably reflect the different lengths of the legs on chairs and sofas, the different distances between their legs, and the different volumes above them. Wu and Barsalou (1999) show that emergent properties arise during productive conceptualization as a result of perceptual simulation.¹⁸

3.1.4. Linguistic control of productivity. A foundational principle in Langacker’s (1986; 1987; 1991; 1997) theory of language is that grammar corresponds to conceptual structure. One dimension of this correspondence is that the productive nature of grammar corresponds to the productive nature of conceptualization. The productive combination of adjectives, nouns, verbs, and other linguistic elements corresponds to the productive combination of perceptual symbols for properties, entities, processes, and other conceptual elements.

This correspondence provides humans with a powerful ability to control each other’s simulations in the absence of the actual referents (Donald 1991; 1993; Tomasello et al. 1993). Without this productive ability, people could only refer to mutually known referents associated with nonproductive linguistic signs. In contrast, the productive ability to construct simulations through language allows people to induce shared simulations of nonexperienced entities and events. Past events experienced by one person can be conveyed to a hearer who has not experienced them, thereby extending the hearer’s realm of knowledge indirectly. Future events can be explored together during planning, decision making, and problem solving as groups of individuals converge on solutions. Because groups can discuss past and future events, greater teamwork becomes possible as does more extensive evaluation of possibilities. The productive control of conceptualization through language appears central to defining what is uniquely human.

3.2. Propositions

Another important lesson that we have learned from amodal symbol systems is that a viable theory of knowledge must implement propositions that describe and interpret situations (e.g., Anderson & Bower 1973; Goodman 1976; Kintsch 1974; Norman et al. 1975; Pylyshyn 1973; 1978; 1981; 1984). A given situation is capable of being construed in an infinite number of ways by an infinite number of propositions. Imagine being in a grocery store. There are limitless ways to describe what is present, including (in amodal form):

CONTAINS (*grocery store, apples*) (3)
ABOVE (*ceiling, floor*)

As these examples illustrate, different construals of the situation result from selecting different aspects of the situation and representing them in propositions. Because an infinite number of aspects can be propositionalized, selecting the propositions to represent a situation is an act of creativity (Barsalou & Prinz 1997). Different construals can also result from construing the same aspects of a situation in different ways, as in:

ABOVE (*ceiling, floor*) (4)
BELOW (*floor, ceiling*)

Bringing different concepts to bear on the same aspects of a situation extends the creative construction of propositions further.

Construals of situations can be arbitrarily complex, resulting from the ability to embed propositions hierarchically, as in:

CAUSE (*HUNGRY* (*shopper*), *BUY* (*shopper, groceries*)) (5)

The productive properties of amodal symbols are central to constructing complex propositions.

Not all construals of a situation are true. When a construal fails to describe a situation accurately, it constitutes a false proposition, as in:

CONTAINS (*grocery store, mountains*) (6)

Similarly, true and false propositions can be negative, as in the true proposition:

NOT (*CONTAINS* (*grocery store, mountains*)) (7)

Thus, propositions can construe situations falsely, and they can indicate negative states.

Finally, propositions represent the gist of comprehension. Comprehenders forget the surface forms of sentences rapidly but remember the conceptual gist for a long time (e.g., Sachs 1967; 1974). Soon after hearing “Marshall gave Rick a watch,” listeners would probably be unable to specify whether they had heard this sentence as opposed to “Rick received a watch from Marshall.” However, listeners would correctly remember that it was Marshall who gave Rick the watch and not vice versa, because they had stored the proposition:

GIVE (*Agent* = marshall, *Recipient* = rick, *Object* = watch) (8)

Thus, propositions capture conceptualizations that can be paraphrased in many ways.

Basically, propositions involve bringing knowledge to bear on perception, establishing type-token relations between concepts in knowledge and individuals in the perceived world. This requires a conceptual system that can

combine types (concepts) productively to form hierarchical structures, and that can then map these structures onto individuals in the world. It is fair to say that this ability is not usually recognized as possible in perceptual theories of knowledge, again because they are widely construed as recording systems. Indeed, this belief is so widespread that the term “propositional” is reserved solely for nonperceptual theories of knowledge. As we shall see, however, if one adopts the core properties of perceptual symbol systems, the important properties of propositions follow naturally. Because perceptual symbol systems have the same potential to implement propositions, they too are propositional systems.¹⁹

3.2.1. Type-token mappings. To see how perceptual symbol systems implement type-token mappings, consider Figure 5A. The large panel with a thick solid border stands for a perceived scene that contains several individual entities. On the far left, the schematic drawing of a jet in a thin solid border stands for the simulator that underlies the concept *jet*. The other schematic drawing of a jet in a thick dashed border represents a specific simulation that provides a good fit of the perceived jet in the scene. Again, such drawings are theoretical notations that should not be viewed as literal images. The line from the simulator to the simulation stands for producing the simulation from the simulator. The line from the simulation to the perceived individual stands for fusing the simulation with the individual in perception.

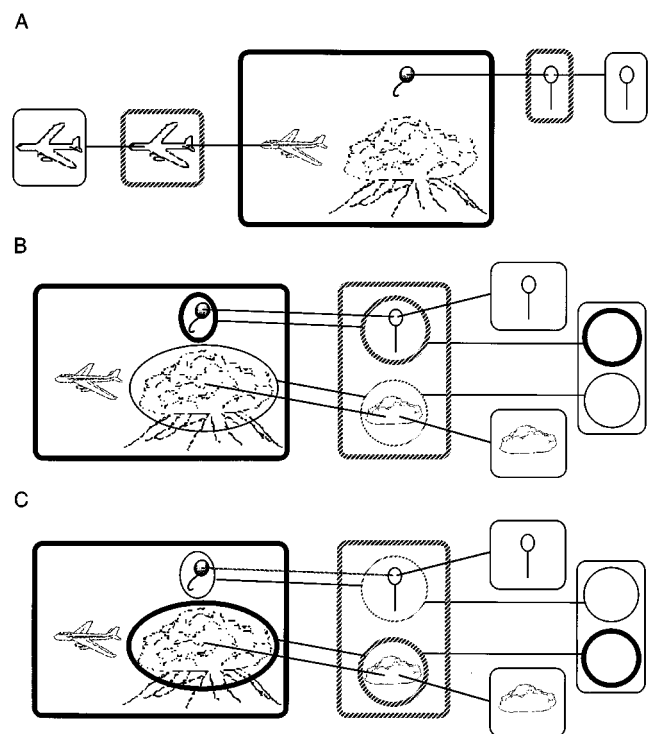


Figure 5. (A) Examples of how perceptual symbol systems represent true propositions by fusing perceived entities with simulations constructed from simulators. (B) Example of a complex hierarchical proposition. (C) Example of an alternative interpretation of the same aspects of the scene in panel B. Boxes with thin solid lines represent simulators; boxes with thick dashed lines represent simulations; boxes with thick solid lines represent perceived situations.

The activation of the simulator for *jet* in Figure 5A results from attending to the leftmost individual in the scene. As visual information is picked up from the individual, it projects in parallel onto simulators in memory. A simulator becomes increasingly active if (1) its frame contains an existing simulation of the individual, or if (2) it can produce a novel simulation that provides a good fit (sect. 2.4.4). The simulation that best fits the individual eventually controls processing.²⁰ Because the simulation and the perception are represented in a common perceptual system, the final representation is a fusion of the two. Rather than being a pure bottom-up representation, perception of the individual includes top-down information from the simulation (sect. 2.4.6). The constructs of *meshing*, *blending*, *perceptual hypotheses*, and *perceptual anticipation* are closely related (Fauconnier & Turner 1998; Glenberg 1997; Kosslyn 1994; Neisser 1976).

Binding a simulator successfully with a perceived individual via a simulation constitutes a type-token mapping. The simulator is a type that construes a perceived token as having the properties associated with the type. Most importantly, this type-token mapping implicitly constitutes a proposition, namely, the one that underlies “It is true that the perceived individual is a jet.” In this manner, perceptual symbol systems establish simple propositions.

3.2.2. Categorical inferences. Binding a simulator to a perceived individual allows perceivers to draw wide variety of categorical inferences (sects. 2.4.4, 2.4.7). If the jet in Figure 5A should become occluded, or if the perceivers should turn away, a simulation of its trajectory could predict where it will reappear later. Simulating the jet might further suggest many aspects of it that are not visible. For example, the simulation might suggest that the jet contains pilots, passengers, luggage, and fuel. Similarly, it might suggest that the jet is likely to fly horizontally, not vertically, and is eventually likely to decrease altitude, put down its wheels, and land at an airport. From bringing the multimodal simulator for *jet* to bear on the perceived individual, a wealth of top-down inference becomes available. Finally, the binding process updates the simulator for *jet*. As selective attention extracts perceptual symbols from the current individual and uses them to construct a simulation, they become integrated into the underlying frame that helped produce it (as illustrated in Fig. 3).

3.2.3. False and negative propositions. So far, we have only considered true propositions, namely, successful bindings between simulators and perceived individuals. Not all attempted bindings succeed, however, and when they do not, they constitute false propositions. Similarly, negative propositions occur when the explicitly noted absence of something in a simulation corresponds to an analogous absence in a perceived situation. Because false and negative propositions receive detailed treatment in section 3.4 on abstract concepts, further discussion is deferred until then.

3.2.4. Multiple construals. Because a given situation can be construed in infinite ways, propositional construal is creative. One way of producing infinite propositions is by selecting different aspects of the situation to construe. Figure 5A illustrates how perceptual symbol systems implement this ability. If perceivers were to focus attention on the uppermost individual in the scene, rather than on the leftmost

individual, they would construe the scene differently. Rather than bringing the simulator for *jet* to bear, they would bring the simulator for *balloon* to bear. The successful mapping that results represents a different proposition, namely, the one that underlies “It is true that the perceived individual is a balloon.” Because infinitely many aspects of the scene can be selected and bound to simulators, an infinite number of propositions describing the scene are possible.

3.2.5. Productively produced hierarchical propositions.

Perceptual symbol systems readily implement complex hierarchical propositions. In Figure 5B, a hierarchical simulation of a balloon above a cloud is constructed productively from the simulators for *balloon*, *cloud*, and *above* (sect. 3.1). This hierarchical simulation in turn is fused successfully with individuals in the perceived scene and the regions they occupy. The result is the representation of the hierarchical proposition that underlies, “It is true that there is a balloon above a cloud.”

3.2.6. Alternative interpretations. Perceptual symbol systems also readily represent alternative interpretations of the same individuals in a scene. As Figure 5C illustrates, the simulator for *below* can be mapped into the same aspects of the perceived situation as the simulator for *above* (Fig. 5B). Because the same spatial configuration of regions satisfies both, either can represent a true proposition about the scene, differing only in where attention is focused (i.e., the upper or lower region). In this manner, and in many others as well, perceptual symbols support different interpretations of the same information. To the extent that different simulations can be fit successfully to the same perceived information, different interpretations result.

3.2.7. Gist. A perceptual simulation represents a gist that can be paraphrased in multiple ways. Imagine that someone hears the sentence, “The balloon is above the cloud.” To represent the sentence’s meaning, the listener might construct a simulation that focuses attention on the upper region of *above*, as in Figure 5B. When later trying to remember the sentence, however, the listener might construct a simulation that has lost information about where attention resided. As a result, it is impossible to specify whether the earlier sentence had been “The balloon is above the cloud” or “The cloud is below the balloon.” As information becomes lost from the memory of a simulation, paraphrases become increasingly likely. Furthermore, because different simulators can often be mapped into the remaining information, additional paraphrases become possible.

In summary, perceptual symbol systems readily implement all the fundamental properties of propositions reviewed earlier for amodal symbols systems. Perceptual symbol systems produce type-token mappings that provide a wealth of inferences about construed individuals. They produce alternative construals of the same scene, either by selecting different aspects of the scene to simulate, or by simulating the same aspects in different ways. They represent complex hierarchical propositions by constructing simulations productively and then mapping these complex simulations into scenes. They represent the gist of sentences with simulations that can be paraphrased later with different utterances. Finally, as described later, they represent false and negative propositions through failed and absent

mappings between simulations and scenes. For all of these reasons, perceptual symbol systems *are* propositional systems.

3.2.8. Intentionality. Propositions exhibit intentionality, namely, the ability of a mental state to be about something (e.g., K. Lehrer 1989). In Figure 5A, for example, the simulation for *jet* has intentionality, because it is about the specific jet in the scene. A major problem for any theory of representation is specifying how mental states come to be about something (or nothing at all). In particular, perceptually based theories of knowledge are often criticized because they seem to imply that the *content* of an image determines its intentionality (or referent). Thus, if someone simulates a jet, the simulation is assumed to be about a jet just like the one simulated.

The content of a symbol does not specify its intentionality, being neither necessary nor sufficient for establishing reference (e.g., Goodman 1976; also see Bloom & Markson 1998). For example, constructing an exact replica of the Empire State Building is not sufficient for establishing reference from the replica to the building (Schwartz 1981). If this replica were placed on a map of the United States at the location of New York City, it could stand for the entire city of New York, not just for the Empire State Building. Similarly, it is not necessary for a symbol that stands for the Empire State Building to bear any similarity to its referent. If the City of New York had a unique, randomly selected number for each of its buildings, the one for the Empire State Building would constitute an arbitrary symbol for it. As these examples illustrate, the degree to which a symbol's content resembles its referent is neither sufficient nor necessary for establishing reference.

There is good reason to believe that perceptual representations can and do have intentionality. Pictures, physical replicas, and movies often refer clearly to specific entities and events in the world (e.g., Goodman 1976; Price 1953). Just because the content of a perceptual representation is not the only factor in establishing reference, it does not follow that a perceptual representation cannot have reference and thereby not function symbolically.

It is also widely believed that factors external to a symbol's content play important roles in establishing its intentionality (e.g., Dretske 1995). A wide variety of external factors is involved, many of them well documented in the literature. For example, definite descriptions often help establish the reference of a symbol (Russell 1919a), as in:

“The computer on my office desk is broken.” (9)

On hearing this sentence, imagine that a listener constructs a perceptual simulation of a computer. The content of this simulation is not sufficient to specify its reference, given that it could refer to many computers. However, when conjoined productively with the simulation that represents the region “on my office desk,” the complex simulation that results establishes reference successfully, assuming that only one computer resides in this region. As this example illustrates, definite descriptions offer one important type of mechanism external to a symbol's content that helps establish its reference, but they are by no means the only way of accomplishing this. Many other important mechanisms exist as well. However, definite descriptions illustrate the importance of external relations that go beyond a symbol's content in establishing its intentionality.

External mechanisms such as definite descriptions can establish reference for either perceptual or amodal symbols (Goodman 1976); because both types of symbols require such mechanisms, neither has an advantage in this regard. Where perceptual symbols do have an advantage is in the ability of their content to play a *heuristic* role in establishing reference. Although perceptual content is rarely definitive for intentionality, it may often provide a major source of constraint and assistance in determining what a symbol is about. In contrast, because the content of an amodal symbol bears no resemblance to its referents, it cannot provide any such guidance.

To see how the content of a perceptual symbol can play a heuristic role in establishing reference, consider again the example in sentence (9). As we saw, the perceptual content of a simulated computer is not sufficient for establishing the speaker's reference, because it could be mapped to many potential computers in the world, as well as to many other entities. Instead, the simulation that represents “on my office desk” was required to fix the reference. Importantly, however, simulating the top of the desk alone is not sufficient to establish reference if there are other things on the desk besides the computer. Although this simulation restricts the range of potential reference to entities on the speaker's desk, it does not specify a unique referent among them. However, the high resemblance of a simulated computer to one of these entities finalizes the act of reference, specifying the single most similar individual, namely, the computer. In this manner, the content of the perceptual symbol for *computer* plays an important – although by no means complete – role in establishing reference. Both components of the definite description – simulating the top of the speaker's desk *and* simulating a computer – are required to refer successfully (Barsalou et al. 1993).²¹

The perceptual symbol that comes to stand for a referent establishes sense as well as reference (Frege 1892/1952). As we saw in Figures 5B and 5C, two different propositions can have the same reference but different interpretations or senses (i.e., a balloon above a cloud versus a cloud below a balloon). Because a perceptual symbol contains content, this content represents the symbol's referents accordingly. As different content becomes active to represent the same referent, it becomes fused with the perceived referent to establish different interpretations (sects. 3.2.2, 3.2.4).²²

3.2.9. The construal of perceptual symbols. The psychological and philosophical literatures raise important problems about the construal of images. These problems do not concern propositional construal per se, because they do not bear directly on the problem of mapping types to tokens. Instead, they concern the construal of images in the absence of reference.

One such problem is the ambiguity of images (Chambers & Reisberg 1992). Consider a mental image of a circular shape. How do we know whether this image represents a pie, a wheel, or a coin? One solution is to assume that a circular shape would rarely, if ever, be imaged in isolation, but would instead be part of a simulation constructed from a simulator (sect. 2.4). Approximately the same schematic circular shape might be produced during the simulation of a pie, a wheel, or a coin. As discussed for background-dependent meaning (sect. 2.5.3), however, this shape would be framed in the context of the particular simulator that produced it, thereby giving it a clear interpretation. Al-

though the shape might be ambiguous in isolation, when linked historically to a particular simulator over the course of a simulation, it is not.

Conversely, another problem of image construal concerns how radically different images can be construed as images of the same thing. Consider Schwartz's (1981) example of a circle. Normally, a circle in full view looks round. If its lower half is occluded, however, it looks like a semi-circle. If partially occluded by visual clutter, it looks like a circular arrangement of line segments. If rotated 45° around its vertical axis, it looks like an oval. If rotated 90°, it looks like a straight vertical line. How do we know that these five different images all represent the same entity? Again, the construct of a simulator provides a solution. If the simulator for *circle* constructs an instance and then transforms it into different surface images, the transformational history of this particular circle links all these images so that they are construed as the same individual. As described for concept stability (sect. 2.4.5), a simulator unifies all the simulations it produces into a single concept. The example here is a similar but more specific case where a simulator unifies all the perspectives and conditions under which a given individual can be simulated.

3.3. Variable embodiment

Variable embodiment is the idea that a symbol's meaning reflects the physical system in which it is represented (Clark 1997; Damasio 1994; Glenberg 1997; Johnson 1987; Lakoff 1987; Lakoff & Johnson 1980; Newton 1996). As different intelligent systems vary physically, their symbol systems are likely to vary as well. Unlike productivity and propositional construal, variable embodiment does not exist in amodal symbol systems. It exists only in perceptual symbol systems, arising from their first two core properties: neural representations in perceptual systems (sect. 2.1), and schematic symbol formation (sect. 2.2). Before considering variable embodiment in perceptual symbol systems, it is useful to consider its absence in amodal symbol systems.

According to the functionalist perspective that has dominated modern cognitive science, the symbol system underlying human intelligence can be disembodied. Once we characterize the computational properties of this system successfully, we can implement it in other physical systems such as computers. Thus, functionalism implies that the computational system underlying human intelligence can be understood independently of the human body. It further implies that the same basic symbol system operates in all normal humans, independent of their biological idiosyncrasies. Just as computer software can be characterized independently of the particular hardware that implements it, the human symbol system can be characterized independently of the biological system that embodies it. For accounts of this view, see Putnam (1960), Fodor (1975), and Pylyshyn (1984); for critiques, see Churchland (1986) and G. Edelman (1992).

The discrete referential nature of amodal symbols lies at the heart of modern functionalism. To see this, consider the transformations of the word "CUP" in Figure 6 and their lack of implications for reference. As Figure 6 illustrates, the word "CUP" can increase in size, it can rotate 45° counterclockwise, and the \supset on the P can become separated from the |. In each case, the transformation implies nothing new about its referent. If "CUP" originally referred to the

referent on the left of Figure 6, its reference does not change across these transformations. Making "CUP" larger does not mean that its referent now appears larger. Rotating "CUP" 45° counterclockwise does not imply that its referent is tipped. Separating the \supset on the P horizontally from the | does not imply that the handle of the cup has now broken off. Because words bear no structural relations to their referents, structural changes in a word have no implications for analogous changes in reference. As long as the conventional link between a word and its referent remains intact, the word refers to the referent discretely in exactly the same way across transformations.

Because amodal symbols refer in essentially the same way as words do, they also refer discretely. Changes in their form have no implications for their meaning. It is this discrete property of amodal symbols that makes functionalism possible. Regardless of how an amodal symbol is realized physically, it serves the same computational function. Physical variability in the form of the symbol is irrelevant, as long as it maintains the same conventional links to the world, and the same syntactic relations to other amodal symbols.

Perceptual symbols differ fundamentally. Unlike an amodal symbol, variations in the form of a perceptual symbol can have semantic implications (cf. Goodman 1976). As Figure 6 illustrates, the schematic perceptual symbol for a cup can increase in size, it can rotate 45° counterclockwise, and the handle can become separated from the cup. In each case, the transformation implies a change in the referent (Fig. 6). Increasing the size of the perceptual symbol implies that the referent appears larger, perhaps because the perceiver has moved closer to it. Rotating the perceptual symbol 45° counterclockwise implies that the cup has tipped analogously. Separating the handle from the cup in the perceptual symbol implies that the handle has become detached from the referent. Because perceptual symbols bear structural relations to their referents, structural changes in a symbol imply structural changes in its referent, at least under many conditions.²³

The analogically referring nature of perceptual symbols makes their embodiment critical. As the content of a symbol varies, its reference may vary as well. If different intelligent systems have different perceptual systems, the conceptual systems that develop through the schematic symbol formation process may also differ. Because their symbols contain different perceptual content, they may refer to different structure in the world and may not be functionally equivalent.

3.3.1. Adaptive roles of variable embodiment. Variable embodiment allows individuals to adapt the perceptual symbols in their conceptual systems to specific environments. Imagine that different individuals consume somewhat different varieties of the same plants because they live in different locales. Through perceiving their respective foods, different individuals develop somewhat different perceptual symbols to represent them. As a result, somewhat different conceptual systems develop through the schematic symbol formation process, each tuned optimally to its typical referents.²⁴

Variable embodiment performs a second useful function, ensuring that different individuals match perceptual symbols optimally to perception during categorization (sect. 2.4.4). Consider the perception of color. Different individuals from the same culture differ in the detailed psy-




Figure 6. An example of how transforming a word or an amodal symbol fails to produce an analogous transformation in reference, whereas transforming a perceptual simulation does.

chophysical structure of their color categories (Shevell & He 1995; V. Smith et al. 1976). Even individuals with normal color vision discriminate the same colors in somewhat different ways, because of subtle differences in their perceptual systems. As a result, when the schematic symbol formation process establishes the ability to simulate colors, somewhat different simulators arise in different individuals. Because each simulator is grounded in a subtly different implementation of the same perceptual system, it represents the same symbols in subtly different manners. Most important, each individual's simulator for color optimally matches their symbols for colors to their perceptions of colors. As different individuals search for bananas at the grocery store, they can simulate the particular representation of *yellow* that each is likely to perceive on encountering physical instances.

Because humans vary on all phenotypic traits to some extent, they are likely to vary on all the perceptual discriminations that could be extracted to form perceptual symbols, not just color. Similar variability should arise in the perception of shape, texture, space, pitch, taste, smell, movement, introspection, and so forth. If so, then variable embodiment allows the human conceptual system to adapt itself naturally to variability in perceptual systems. In contrast, such adaptability is not in the spirit of functionalism. Because functionalism rests on amodal symbols that bear no structural relation to their referents, it neither anticipates nor accommodates individual variability in perception.

3.3.2. Variable embodiment in conceptual variability and stability.

Earlier we saw that conceptual variability arises out of simulators (sect. 2.4.5). When the simulator for a category produces different simulations, conceptual variability arises both between and within individuals. Variable embodiment provides a second source of variability. Different individuals represent the same concept differently because their perceptual symbols become tuned to somewhat different physical environments and develop in somewhat different perceptual systems.

We also saw earlier that simulators provide conceptual stability (sects. 2.4.5, 3.2.9). When different simulations can be traced back to the same simulator, they become unified as instances of the same concept. Embodiment similarly provides stability. Although perceptual systems produce variable embodiment across different individuals, they also produce *shared embodiment* at a more general level. Because most humans have roughly the same mechanisms for perceiving color, they have roughly the same conceptual representations of it. Although perceptual systems induce idiosyncrasies in perception and conception, they also induce basic commonalities (Newton 1996).

3.4. Abstract concepts

Representing abstract concepts poses a classic challenge for perceptual theories of knowledge. Although representing concrete objects has sometimes seemed feasible, repre-

senting events, mental states, social institutions, and other abstract concepts has often seemed problematic.

3.4.1. Representing abstract concepts metaphorically.

Cognitive linguists have suggested that metaphor provides a perceptually based solution to the representation of abstract concepts (e.g., Gibbs 1994; Johnson 1987; Lakoff 1987; Lakoff & Johnson 1980; Lakoff & Turner 1989; Turner 1996). For example, Lakoff and Johnson (1980) suggest that the concrete domain of *liquid exploding from a container* represents the abstract concept of *anger*. Although metaphor most certainly plays a major role in elaborating and construing abstract concepts, it is not sufficient for representing them (Barsalou et al. 1993; Murphy 1996). Instead, a direct, nonmetaphorical representation of an abstract domain is essential for two reasons: first, it constitutes the most basic understanding of the domain. Knowing only that *anger* is like *liquid exploding from a container* hardly constitutes an adequate concept. If this is all that people know, they are far from having an adequate understanding of *anger*. Second, a direct representation of an abstract domain is necessary to guide the mapping of a concrete domain into it. A concrete domain cannot be mapped systematically into an abstract domain that has no content.

As research on emotion shows (e.g., G. Mandler 1975; Shaver et al. 1987; Stein & Levine 1990), people have direct knowledge about *anger* that arises from three sources of experience. First, *anger* involves the appraisal of an initiating event, specifically, the perception that an agent's goal has been blocked. Second, *anger* involves the experience of intense affective states. Third, *anger* often involves behavioral responses, such as expressing disapproval, seeking revenge, and removing an obstacle. As people experience each of these three aspects of *anger*, they develop knowledge directly through the schematic symbol formation process (sect. 2.2).

Although people may understand *anger* metaphorically at times, such understanding elaborates and extends the direct concept. Furthermore, metaphorical language may often indicate polysemy rather than metaphorical conceptualization (Barsalou et al. 1993). For example, when someone says, "John exploded in anger," the word "explode" may function polysemously. "Explode" may have one sense associated with heated liquid exploding in containers, and another associated with the rapid onset of angry behavior. Rather than activating conceptual knowledge for *liquid exploding from a container*, "explode" may simply activate a perceptual simulation of angry behavior directly. As in the direct interpretation of indirect speech acts that bypass literal interpretation (Gibbs 1983; 1994), familiar metaphors may produce direct interpretations that bypass metaphorical mappings. Just as "Can you pass the salt?" bypasses its literal meaning to arrive directly at a pragmatic request, so can "explode" bypass its concrete sense to arrive directly at its introspective sense. Although novel metaphors may typically require a metaphorical mapping, familiar metaphors may bypass this process through polysemy.

3.4.2. Representing abstract concepts directly with perceptual symbols.

Ideally, it should be shown that perceptual symbol systems can represent all abstract concepts directly. Such an analysis is not feasible, however, given the large number of abstract concepts. An alternative strategy

is to select quintessential abstract concepts and show that perceptual accounts are possible. The next two subsections provide perceptual accounts of two such concepts, *truth* and *disjunction*, as well as related concepts in their semantic fields. If challenging abstract concepts like these can be represented perceptually, we have good reason to believe that other abstract concepts are also tractable.

In developing perceptual accounts of these abstract concepts and others, a general trend has emerged. Across these concepts, three mechanisms appear central to their representation. First, an abstract concept is framed against the background of a simulated event sequence. Rather than being represented out of context in a single time slice, an abstract concept is represented in the context of a larger body of temporally extended knowledge (Barwise & Perry 1983; Fillmore 1985; Langacker 1986; Newton 1996; Yeh & Barsalou 1999a; 1999b). As we saw earlier in the sections on simulators (sect. 2.4) and framing (sect. 2.5.3), it is possible to simulate event sequences perceptually, and it is possible for a simulator to frame more specific concepts. Thus, perceptual symbol systems can implement the framing that underlies abstract concepts.

Second, selective attention highlights the core content of an abstract concept against its event background (Langacker 1986; Talmy 1983). An abstract concept is not the entire event simulation that frames it but is a focal part of it. As we saw earlier in the section on schematic symbol formation (sect. 2.2), it is possible to focus attention on a part of a perceptual simulation analytically. In this way, perceptual symbol systems capture the focusing that underlies abstract concepts.

Third, perceptual symbols for introspective states are central to the representation of abstract concepts. If one limits perceptual symbols to those that are extracted from perception of the external world, the representation of abstract concepts is impossible. As we saw earlier in the section on multimodal symbols, the same symbol formation process that operates on the physical world can also operate on introspective and proprioceptive events. As a result, the introspective symbols essential to many abstract concepts can be represented in a perceptual symbol system. Although many different introspective events enter into the representation of abstract concepts, propositional construal appears particularly important (sect. 3.2). As we shall see, the act of using a mental state to construe the physical world is often central.

Together, these three mechanisms – framing, selectivity, and introspective symbols – provide powerful tools for representing abstract concepts in perceptual symbol systems. As we shall see shortly, they make it possible to formulate accounts of *truth*, *disjunction*, and related concepts. This success suggests a conjecture about the representation of abstract concepts: framing, selectivity, and introspective symbols allow a perceptual symbol system to represent any abstract concept. This conjecture in turn suggests a general strategy for discovering these representations. First, identify an event sequence that frames the abstract concept. Second, characterize the multimodal symbols that represent not only the physical events in the sequence but also the introspective and proprioceptive events. Third, identify the focal elements of the simulation that constitute the core representation of the abstract concept against the event background. Finally, repeat the above process for any other event sequences that may be relevant to representing the

concept (abstract concepts often refer to multiple events, such as *marriage* referring to a ceremony, interpersonal relations, domestic activities, etc.). If the conjecture is correct, this strategy should always produce a satisfactory account of an abstract concept.

Of course, the success of this strategy does not entail that people actually represent abstract concepts this way. Such conclusions await empirical assessment and support. However, success is important, because it demonstrates that perceptual symbol systems have the expressive power to represent abstract concepts. In a pilot study with Karen Olseth Solomon, we have obtained preliminary support for this account. Subjects produced more features about event sequences and introspection for abstract concepts, such as *truth* and *magic*, than for concrete concepts, such as *car* and *watermelon*. Recent neuroscience findings are also consistent with this proposal. As reviewed by Pulvermüller (1999), abstract concepts tend to activate frontal brain regions. Of the many functions proposed for frontal regions, two include the coordination of multimodal information and sequential processing over time. As just described, abstract concepts exhibit both properties. On the one hand, they represent complex configurations of multimodal information that must be coordinated. On the other, they represent event sequences extended in time. Thus, frontal activation is consistent with this proposal.

3.4.3. Representing *truth*, *falsity*, *negation*, and *anger*.

This analysis does not attempt to account for all senses of *truth*, nor for its formal senses. Only a core sense of people's intuitive concept of *truth* is addressed to illustrate this approach to representing abstract concepts. In the pilot study with Solomon, subjects described this sense frequently. Figure 7A depicts the simulated event sequence that underlies it. In the first subevent, an agent constructs a perceptual simulation of a balloon above a cloud using the productive mechanisms described earlier. The agent might have constructed this simulation on hearing a speaker say, "There's a balloon above a cloud outside," or because the agent had seen a balloon above a cloud earlier, or for some other reason. Regardless, at this point in the event sequence, the agent has constructed a simulation. In the second subevent, the agent perceives a physical situation (i.e., the scene inside the thick solid border), and attempts to map the perceptual simulation into it. The agent may attempt this mapping because a speaker purported that the statement producing the simulation was about this particular situation, because the agent remembered that the simulation resulted from perceiving the situation earlier, or for some other reason. The agent then assesses whether the simulation provides an accurate representation of the situation, as described earlier for propositional construal. In this case, it does. Analogous to the simulation, the situation contains a balloon and a cloud, and the balloon is above the cloud. On establishing this successful mapping, the agent might say, "It's true that a balloon is above a cloud," with "true" being grounded in the mapping.

This account of *truth* illustrates the three critical mechanisms for representing abstract concepts. First, a simulated event sequence frames the concept. Second, the concept is not the entire simulation but a focal part of it, specifically, the outcome that the simulation provides an accurate construal of the situation. Third, perceptual symbols for introspective events are central to the concept, includ-




Figure 7. (A) Accounting for one sense of *truth* using perceptual symbols. (B) Accounting for one sense of *falsity* using perceptual symbols. (C) Accounting for one sense of *negation* using perceptual symbols. Boxes with thin solid lines represent simulations; boxes with thick dashed lines represent simulations; boxes with thick solid lines represent perceived situations.

ing those for a perceptual simulation, the process of comparing the perceptual simulation to the perceived situation, and the outcome of establishing a successful mapping between them. After performing this complex event sequence on many occasions, a simulator develops for *truth*, that is, people learn to simulate the experience of successfully mapping an internal simulation into a perceived scene.²⁵

The concept of *falsity* is closely related to the concept of *truth*. The account here addresses only one sense of people's intuitive concept for *falsity*, although it too is polysemous and has formal interpretations. As Figures 7A and 7B illustrate, very similar event sequences underlie these two concepts. In both, a simulation is constructed initially, followed by the perception of a situation. The two sequences differ only in whether the simulation can or cannot be mapped successfully into the situation. Whereas the mapping succeeds for *truth*, it fails for *falsity*. Thus, a speaker, after failing to map the simulation into the situation in Figure 7B, might say, "It is false that there is a balloon above a cloud," with "false" being grounded in the failed mapping. The slant mark through the line between the simulated balloon and its absence in perception is a theoretical device, not a cognitive entity, that stands for a failed attempt at fusing them.²⁶

The concept of *negation* is also closely related to the concept of *truth*. Again, the account here addresses only one sense of people's intuitive concept. As Figure 7C illustrates, negation results from explicitly noting the absence of a binding between a simulator and a simulation. In this particular case, the simulator for *balloon* is not bound to anything in the upper region of the *above* simulation. Explicitly noting the absence of a binding becomes part of the


simulation, as indicated theoretically by the line with a slant mark through it. When the simulation is compared subsequently to the perceived situation, the absent mapping underlies the negated claim, “It is true that there is not a balloon above a cloud.” Negation can also occur with falsity. If the perceived situation in Figure 7C *did* contain a balloon above a cloud, the absent mapping would fail to map into the situation, and the perceiver might say, “It is false that there is not a balloon above a cloud.” As these examples illustrate, falsity and negation can be represented with absent mappings between simulators, simulations, and perceived situations. Givón (1978) and Glenberg et al. (1998a) make related proposals.

The concept of *anger* also belongs to this semantic field. As discussed earlier (sect. 3.4.1), a core component of *anger* is the appraisal of a blocked goal. In perceptual symbol systems, a goal is a simulated state of the world that an agent would like to achieve, and a blocked goal is the failed mapping of such a simulation at a time when it is expected to map successfully. Thus, a blocked goal involves an event sequence very similar to the one in Figure 7B for *falsity*. In both, a simulated situation fails to match a perceived situation. As also described earlier, *anger* involves other core components, such as affective experience and behavioral responses, which are likely to be included in its simulations. Nevertheless, *anger* belongs to the same semantic field as *truth* and *falsity* by virtue of sharing a common event sequence, namely, assessing the fit of a simulation to a situation. *Lie* is also related, where a statement induces a simulation purported to be true that is actually not (i.e., the simulation is negative in the liar’s simulation but false in the deceived’s simulation).

Finally, it is worth noting that the experiential basis for this semantic field probably exists from birth, if not before. Anyone who has raised children knows that an infant has expectations about food, comfort, and sleep. When these expectations are satisfied, the infant is content; when they are not, the infant conveys dissatisfaction. The point is that the infant constantly experiences the event sequence that underlies *truth*, *falsity*, *negation*, *anger*, and their related concepts. From birth, and perhaps before, the infant simulates its expected experience and assesses whether these simulations map successfully into what actually occurs (cf. Haith 1994). As an infant experiences this event sequence day after day, it learns about satisfied versus failed expectations. If infants have the schematic symbol formation process described earlier, and if they can apply it to their introspective experience, they should acquire implicit understandings of *truth*, *falsity*, and *negation* long before they begin acquiring language. Once the language acquisition process begins, a rich experiential basis supports learning the words in this semantic field.

3.4.4. Representing *disjunction* and ad hoc categories.

Disjunction, like *truth*, is polysemous and includes formal senses. The analysis here only attempts to account for one intuitive sense. As Figure 8 illustrates, an event sequence is again critical. In the first subevent, an agent perceives a situation and constructs an internal simulation that construes parts of it propositionally. In the second subevent some time later, the agent attempts to remember the earlier situation, searching for the propositions constructed at that time. As the second subevent illustrates, the agent retrieves one proposition for *between* and two others for *tree* and *barn*,



“There was a horse or a cow between the tree and the barn.”

Figure 8. Accounting for one sense of *disjunction* using perceptual symbols. Boxes with thin solid lines represent simulators; boxes with thick dashed lines represent simulations; the box with the thick solid line represents a perceived situation. The dashed lines in Subevent 3 represent alternating simulations of a horse and a cow in the middle region of the simulated event.

which are combined productively. Although the tree on the left and the barn on the right have been remembered, the *between* proposition implies that there was an entity in the middle, which has been forgotten. Finally, in the third subevent, the agent attempts to reconstruct the entity that could have existed in the middle. Using reconstructive mechanisms (irrelevant to an analysis of *disjunction*), the agent simulates two possible entities in this region, a *horse* and a *cow*, with the intent that they possibly construe the region’s original content. The dotted lines are theoretical notations which indicate that the agent is not simulating these two specializations simultaneously but is alternating between them in a single simulated event, while holding specializations of the outer regions constant. During this process, the agent might say, “There was a horse or a cow between the tree and the barn,” with “or” being grounded in the alternating simulations of the middle region.

Similar to *truth* and *negation*, the three mechanisms for representing abstract concepts are present. Again, an event sequence frames the concept, while attention focuses on the core content, namely, the alternating specializations of the middle region. Again, introspective events, such as productive simulation, propositional construal, and alternating specialization, are central.²⁷

The semantic field that contains *disjunction* also contains ad hoc categories (Barsalou 1983; 1985; 1991). As Barsalou (1991) proposes, ad hoc categories construe entities as playing roles in events. For example, the category of *things to stand on to change a light bulb* construes a chair as playing the role of allowing someone to reach a light bulb while replacing it. Alternatively, the category of *things used to prop a door open* construes a chair as playing the role of keeping a door from closing, perhaps to cool off a hot room. Each of these categories is disjunctive, because a wide variety of entities can play the same role in an event.

Figure 9 illustrates how perceptual symbol systems represent ad hoc categories. Figure 9A stands for simulating the event of changing a light bulb, and Figure 9C stands for simulating the event of propping a door open. In each simulation, a schematic region can be specialized disjunctively, thereby forming an ad hoc category. In Figure 9A, a schematic region exists for the entity that is stood on to

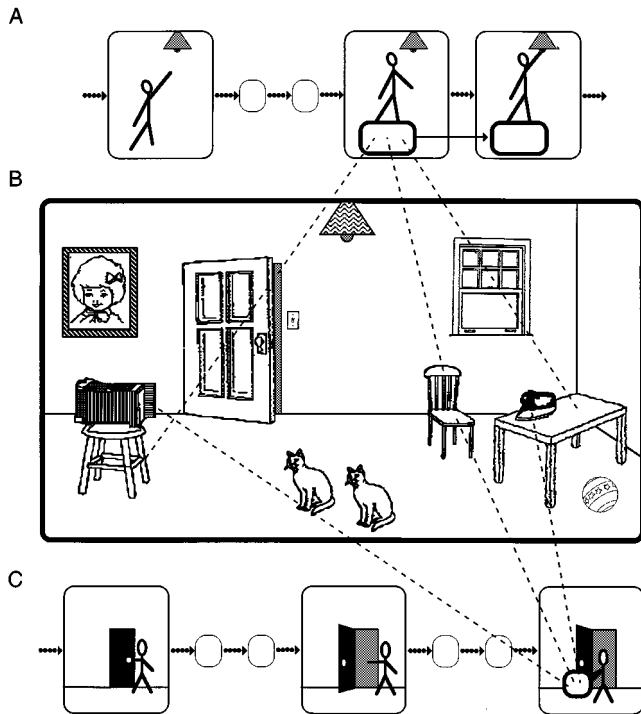


Figure 9. Accounting for the ad hoc categories of *things to stand on to change a light bulb* (A) and *things that could prop a door open* (C), which construe common entities in the same scene differently (B).

reach the light bulb. In Figure 9C, a schematic region exists for the entity that props the door open. Each of these schematic regions can be specialized with entities in the perceived situation (Fig. 9B). The schematic region for *something to stand on to change a light bulb* in Figure 9A can be specialized with the stool, chair, and table, whereas the schematic region for *something used to prop a door open* in Figure 9C can be specialized with the accordion, the chair, and the iron.²⁸

As Figure 9 further illustrates, the same entity is construed in different ways when it specializes regions in different event simulations. For example, the chair in Figure 9B specializes regions in simulations for changing a light bulb and for propping a door open. As in Figures 5B and 5C, a perceptual symbol system represents different interpretations of the same individual by binding it to different simulations. A wide variety of other role concepts, such as *merchandise*, *food*, and *crop*, can similarly be represented as disjunctive sets that instantiate the same region of a simulation.

3.4.5. Summary. Perceptual symbol systems can represent some of the most challenging abstract concepts directly. Although they have not been shown to represent all abstract concepts, their ability to represent some of the most difficult cases is encouraging. Such success suggests that other abstract concepts can be represented similarly by (a) identifying the event sequences that frame them, (b) specifying the physical and introspective events in these sequences, and (c) identifying the focal elements of these simulations. As these examples further illustrate, abstract concepts are perceptual, being grounded in temporally extended simulations of external and internal events. What may make these concepts seem nonperceptual is their heavy reliance

on complex configurations of multimodal information distributed over time. This property distinguishes abstract concepts from more concrete concepts, but it does not mean that they are nonperceptual.

4. Implications

A perceptual theory of knowledge is not restricted to a recording system that can only store and process holistic images. Construing perceptual theories this way fails to address the most powerful instances of the class. As we have seen, it is possible to formulate a perceptual theory of knowledge that constitutes a fully functional conceptual system. According to this theory, selective attention extracts components of perceptual experience to establish simulators that function as concepts. Once established, simulators represent types, they produce categorical inferences, they combine productively to form complex simulations never experienced, and they establish propositions that construe individuals in the world. Furthermore, this framework is not limited to the representation of concrete objects but can also represent abstract concepts. A perceptual symbol system can perform all of the important functions that modern cognitive science asks of a theory of knowledge.

Once one begins viewing human knowledge this way, it changes how one views other aspects of cognition. Thinking about memory, language, and thought as grounded in perceptual simulation is very different from thinking about them as grounded in the comparison of feature lists, the search of semantic nets, or the left-to-right processing of predicates. The remaining sections explore how adopting the perspective of perceptual symbol systems affects one's views of cognition, neuroscience, evolution, development, and artificial intelligence.

4.1. Implications for cognition

4.1.1. Perception. As described earlier, cognition diverged from perception in the twentieth century as knowledge came to be viewed as nonperceptual. Conversely, perception diverged from cognition, focusing primarily on bottom-up sensory mechanisms and ignoring top-down effects. From the perspective of perceptual symbol systems, the dissociation between perception and cognition is artificial and invalid. Because perception and cognition share common neural systems, they function simultaneously in the same mechanisms and cannot be divorced. What is learned about one becomes potentially relevant for understanding the other.

This influence is not unidirectional: cognition does not become more perceptual while perception remains unaffected. Perception is not an entirely modular system with cognition lying outside it. Because perception shares systems with cognition, bottom-up activation of perceptual systems engages cognitive processes immediately. As described earlier (sect. 2.4.6), bottom-up information may dominate conflicting top-down information but fuse with consistent top-down information.

4.1.2. Categorization and concepts. From the perspective of perceptual symbol systems, categorization is propositional construal (sect. 3.2). The perception of an individual activates the best-fitting simulator, thereby establishing a type-token mapping. In the process, the simulator runs a

schematic simulation in the same perceptual systems that perceive the individual, with the simulation adding inferred features. Because the simulation and the perception are represented in shared perceptual systems, they become fused together. In contrast, standard views assume that amodal features in a nonperceptual system become active to represent a perceived individual, with the cognitive and perceptual representations remaining separate.

The concepts that underlie categorization also take on new forms from the perspective of perceptual symbol systems. Rather than being a static amodal structure, a concept is the ability to simulate a kind of thing perceptually. Because simulators partially reproduce perceptual experience, spatiotemporal relations and transformations become central to concepts. Whereas such knowledge is cumbersome and brittle in amodal symbol systems (e.g., Clark 1997; Glasgow 1993; McDermott 1987; Winograd & Flores 1987), it “rides for free” in perceptual symbol systems (Goldstone & Barsalou 1998). Conceptual combination similarly takes on a new character in perceptual symbol systems. Rather than being the set-theoretic combination of amodal features, it becomes the hierarchical construction of complex perceptual simulations.

Perceptual symbol systems also offer a different perspective on the proposal that intuitive theories provide background knowledge about concepts (Murphy & Medin 1985). As we saw earlier (sect. 2.5.3), concepts are often framed in the context of perceptually simulated events. This type of framing may underlie many of the effects attributed to intuitive theories, and it may account for them in a more implicit and less formal manner. Rather than assuming that people use theories to frame concepts, it may be more realistic to assume that people use simulated events from daily experience. For example, an intuitive theory about biology is often assumed to frame concepts of natural kinds. Alternatively, simulations of biological events may frame these concepts and represent implicitly the theoretical principles that underlie them. Specifically, simulations of birth, growth, and mating may frame concepts like *animal* and represent theoretical principles such as *mothers give birth to babies, who grow up to produce their own offspring*. Simulations of various transformations may further support inferences that distinguish natural kinds from artifacts, such as the ability of *leg* to bend for *human* but not for *table*. Similarly, simulations of rolling out pizzas versus the minting of quarters may underlie different inferences about variability in diameter (Rips 1989). Event simulations may also represent functions, such as the function of cups being represented in simulations of drinking from them. Whenever an intuitive theory appears necessary to explain a conceptual inference, a simulated event sequence may provide the requisite knowledge.

4.1.3. Attention. Attention is central to schematic symbol formation. By focusing on an aspect of perceptual experience and transferring it to long-term memory, attention overcomes the limits of a recording system. Attention is a key analytic mechanism in parsing experience into the schematic components that ultimately form concepts. Although attention has played important roles in previous theories of concepts (e.g., Nosofsky 1984; Trabasso & Bower 1968), its role here is to create the schematic perceptual symbols that compose simulators.

Once a simulator becomes established, it in turn controls

attention (cf. Logan 1995). As a simulator produces a simulation, it controls attention across the simulation. Thus, attention becomes a semantic feature, as when it focuses on the upper versus the lower region of the same spatial representation to distinguish *above* from *below*. The control of attention can also contrast a focal concept against a background simulation, as we saw in the sections on framing (sect. 2.5.3) and abstract concepts (sect. 3.4). As these examples illustrate, attention takes on important new roles in the context of perceptual symbol systems.

Perceptual symbol systems also provide natural accounts of traditional attentional phenomena. For example, automatic processing is the running of a highly compiled simulation, whereas strategic processing is the construction of a novel simulation using productive mechanisms. Priming is essentially perceptual anticipation (Glenberg 1997; Neisser 1976), namely, the top-down activation of a simulation that matches a perceptual input and facilitates its processing. In contrast, inhibition is the top-down suppression of a simulator.

4.1.4. Working memory. Current accounts of working memory construe it as a complex set of limited capacity mechanisms, including an executive processor and several modality specific buffers, such as an articulatory loop and a visual short-term memory (Baddeley 1986). From the perspective of perceptual symbol systems, working memory is the system that runs perceptual simulations. The articulatory loop simulates language just heard or about to be spoken. The visual short-term buffer simulates visual experience just seen or currently imagined. The motor short-term buffer simulates movements just performed or about to be performed. The gustatory short-term buffer simulates tastes just experienced or anticipated. The executive processor simulates the execution of procedures just executed or about to be performed. Not only do these working memory systems operate during perception, movement, and problem solving, they can also be used to simulate these activities offline.

Standard theories of cognition assume that working memory contains perceptual representations, whereas long-term memory contains amodal representations (e.g., Baddeley 1986; Kosslyn 1980). From the perspective of perceptual symbol systems, both systems are inherently perceptual, sharing neural systems with perception. Whereas long-term memory contains simulators, working memory implements specific simulations.

4.1.5. Long-term memory. Encoding information into long-term memory is closely related to categorization, because both involve propositional construal. When an individual is encountered, its categorization into a type-token relation encodes a proposition (sect. 3.2.1). To the extent that this proposition receives processing, it becomes established in long-term memory. The many varieties of elaboration and organization in the encoding literature can all be viewed in this manner. Indeed, various findings suggest that elaboration and organization are inherently perceptual (e.g., Brandimonte et al. 1997; Intraub et al. 1998; M. Martin & Jones 1998). From the perspective of perceptual symbol systems, encoding produces a fusion of bottom-up sensation and top-down simulation (sect. 2.4.7). Under conceptually driven orienting tasks, a fusion contains increased information from simulation; under data-driven orienting tasks, a fusion contains increased information from sensation (cf. Jacoby 1983). Furthermore, rich sensory informa-

tion suggests that the memory resulted from perception, not imagination (Johnson & Raye 1981).

Once a type-token fusion becomes stored in long-term memory, it can be retrieved on later occasions. Because the fusion is a perceptual representation, its retrieval is essentially an attempt to simulate the original entity or event. Thus, memory retrieval is another form of perceptual simulation (Glenberg 1997; cf. Conway 1990; Kolers & Roediger 1984), with fluent simulations producing attributions of remembrance (Jacoby et al. 1989). Such simulations can become active unconsciously in implicit memory, or they can become active consciously in explicit memory. Reconstructive memory reflects the unbidden contribution of a simulator into the retrieval process. As a memory is retrieved, it produces a simulation of the earlier event. As the simulation becomes active, it may differ somewhat from the original perception, perhaps because of less bottom-up constraint. To the extent that the remembered event's features have become inaccessible, the simulator converges on its default simulation. The result is the wide variety of reconstructive effects reported in the literature.

4.1.6. Language processing. Language comprehension can be viewed as the construction of a perceptual simulation to represent the meaning of an utterance or text. As suggested earlier, the productive properties of language guide this process, with a combination of words providing instructions for building a perceptual simulation (Langacker 1986; 1987; 1991; 1997; MacWhinney 1998). A typical sentence usually contains reference to specific individuals, as well as predications of them. For example, "the cup on Anna's desk is blue" refers to a particular cup on a particular desk, and predicates that it is blue. To construct a simulation of this sentence, the comprehender simulates the individual desk and cup and then specializes the color of the cup. Later sentences update the simulation by changing the individuals present and/or transforming them. Thus, "it contains pens and pencils" adds new individuals to the simulation, inserting them inside the cup. The affordances of a simulation may often produce inferences during comprehension. For example, spatial properties of the pens, pencils, and cup determine that the pens and pencils sit vertically in the cup, leaning slightly against its lip. If the pens and pencils had instead been placed in a drawer, their orientation would have been horizontal. In an amodal representation, such inferences would not be made, or they would require cumbersome logical formulae.

As individuals and their properties become established in a simulation, they become fused with the simulators that construe them. Thus, the pens, pencils, and cup become fused with simulators for *pen*, *pencil*, and *cup*, and the blue color of the cup becomes fused with *blue*. As a result, the simulators produce inferences about the individuals as needed. For example, if the text stated, "Anna checked whether one of the pens worked, and it didn't," the simulator for *pen* might simulate inferences such as *Anna pressed a button on top of a pen* and *when the pen didn't work, an uncomfortable feeling resulted from the dry pen scraping across the paper*. Similarly, if the text stated, "Anna would have preferred a cup in a lighter shade of blue," the simulator for *blue* might simulate a lighter shade that simulates the cup negatively but that produces a more positive affective response in the simulation of Anna's mental state. As comprehension proceeds, representations of individuals develop, as in the perception of a

physical scene. Simultaneously, simulators become fused with these imagined individuals to represent propositions about them, much like the type-token mappings that develop during the categorization of physical entities.²⁹

As these examples illustrate, perceptual simulation offers a natural account of how people construct the meanings of texts, or what other researchers have called *situation models* and *mental models* (e.g., Johnson-Laird 1983; Just & Carpenter 1987; van Dijk & Kintsch 1983). A variety of findings can be interpreted as evidence that perceptual simulation underlies these models (e.g., Black et al. 1979; Bransford & Johnson 1973; Gernsbacher et al. 1990; Gibbs 1994; Glenberg et al. 1987; Intraub & Hoffman 1992; Morrow et al. 1987; Potter & Faulconer 1975; Potter et al. 1986; Potter et al. 1977; Von Eckardt & Potter 1985; Rinck et al. 1997; Wilson et al. 1993).

4.1.7. Problem solving, decision making, and skill. From the perspective of perceptual symbol systems, problem solving is the process of constructing a perceptual simulation that leads from an initial state to a goal state. Problem solvers can work forward from the initial state or backward from the goal state, but in either case they attempt to simulate a plan that achieves the goal. In novice problem solving, the difficulty is finding satisfactory components of the simulation and ordering them properly. If novices start from the initial state, they may not know the next event to simulate that will lead to the goal. A component may be added to the simulation that has been successful for achieving this type of goal in the past, or a component might be added because its simulated affordances suggest that it may work. For example, one might want to remove caulk between a wall and the counter in a kitchen. If one has never done this before, various plans can be simulated to see which works best (e.g., using a scraper versus a chemical solvent).

Decision making can be viewed as specializing a simulated plan in different ways to see which specialization produces the best outcome (cf. the simulation heuristic of Kahneman & A. Tversky 1982). In evaluating plans to remove caulk from a joint, a decision must be made about how to specialize the region of the handheld instrument. As possible specializations are retrieved, each is simulated to assess which works best. A wide variety of decisions can be viewed this way, including decisions about purchases, occupations, social events, and so forth. For each, an agent simulates a plan to achieve a goal and then tries out disjunctive specializations of a region to see which yields the most promising outcome. In essence, making a decision involves constructing an ad hoc category for a plan region and selecting one of its members (sect. 3.4.4).

Skill results from compiling simulations for most of the plans in a domain through extensive experience (cf. Anderson 1993; Logan 1988; Newell 1990). Rather than searching for plan components and making decisions about how to specialize plan regions, experts retrieve compiled simulations that achieve goals with minimal transformation. During plan execution, simulated plans run slightly ahead of perception. As in propositional construal, simulation and perception become closely entrained. Expertise is achieved when an agent can almost always simulate what is about to occur, rarely stopping to revise the simulation.

4.1.8. Reasoning and formal symbol manipulation. To see how perceptual symbol systems could underlie logical rea-

soning, consider modus ponens. In modus ponens, if the premise $X \rightarrow Y$ is true, and if the premise X is also true, then the conclusion Y follows. Shortly, a formal account of modus ponens in perceptual symbol systems will be presented, but first an implicit psychological account is developed. Imagine that an agent is told, *If a computer is a Macintosh, then it has a mouse*. On hearing this, the agent constructs a perceptual simulation of a Macintosh that includes a mouse, thereby representing the premise, $X \rightarrow Y$, informally (cf. Johnson-Laird 1983). On a later occasion, when a particular Macintosh is encountered (i.e., premise X), it activates the simulation of a Macintosh constructed earlier, which includes a mouse. As the simulation is fused with the individual, a mouse is simulated, even if a physical mouse is not perceived. Through this psychological analogue to modus ponens, the inference that the individual has a mouse is drawn (i.e., Y).

Similar use of simulation could underlie syllogisms such as *Every B is C, A is B, therefore A is C*. Imagine that, over time, an agent experiences a mouse with every Macintosh. As a result, mice become strongly established in the simulator for Macintosh, so that all simulations of a Macintosh include one (i.e., Every B is C). Later, when a new entity is categorized as a Macintosh (i.e., A is B), the simulator for Macintosh produces a simulation that includes a mouse, thereby drawing the syllogistic inference (i.e., A is C).

To the extent that C does not always covary with B, the certainty of C decreases, giving the inference a statistical character (Oaksford & Chater 1994). If an agent has experienced Macintoshes without mice, the inference that a perceived individual has one is less certain, because simulations can be constructed without them as well as with them. To the extent that simulations with mice are easier to construct than simulations without mice, however, the inference is compelling. As a simulation becomes more fluent, its perceived likelihood increases (Jacoby et al. 1989).

Widespread content effects in reasoning are consistent with perceptual simulation. As many researchers have reported, reasoning improves when the abstract variables in arguments are replaced with familiar situations. For example, people draw the invalid inference of affirming the consequent in arguments stated with abstract variables (i.e., receiving the premises $X \rightarrow Y$ and Y , and then concluding X). However, in a familiar domain such as computers, if X is a Macintosh and Y is a mouse, people do not affirm the consequent, because they can think of non-Macintosh computers that have mice. From the perspective of perceptual symbol systems, this improvement reflects the ability to construct relevant simulations (cf. Johnson-Laird 1983). To the extent that the critical events have been experienced, information becomes stored that produces the simulations necessary to drawing only valid inferences.

Perceptual simulation offers a similar account of causal reasoning. Cheng and Novick (1992) propose that the strength of a causal inference reflects the difference between the probability of an event leading to an outcome and the probability of the event not leading to the outcome, with increasingly large differences producing increasingly strong inferences. A perceptual symbol system can compute these differences using perceptual simulations. To the extent that it is easier to simulate an event leading to an outcome than the event not leading to the outcome, the event is construed as a likely cause of the outcome. Indeed, people appear to construct such simulations to assess causality

(Ahn & Bailenson 1996; Ahn et al. 1995). Covariation is also critical, though, because the more the event and outcome covary, the greater the fluency of the simulation, and the stronger the causal attribution (cf. Jacoby et al. 1989). Much additional research in social cognition illustrates that the ease of simulating a scenario underlies the acceptability of a causal explanation (e.g., K. Markman et al. 1993; Pennington & Hastie 1992; Wells & Gavinski 1989).

Finally, it is possible to explain formal symbol manipulation in logic and mathematics through the simulation of *arbitrary* symbols. From perceptual experience with external symbols and operations, the ability to construct analogous simulations internally develops. For example, after watching an instructor work through modus ponens, a student may develop the ability to simulate the formal procedure internally. The student learns to simulate the two premises followed by the inferred conclusion, analogous to how they would be manipulated externally (i.e., simulate " $X \rightarrow Y$ " and " X " as given, then simulate " Y " as true). Furthermore, the student develops the ability to simulate, replacing variables with constants. Thus, if students receive "Macintosh \rightarrow mouse, Macintosh" in a problem, they know that memory should be searched for a simulated rule that this pattern of constants can specialize. On retrieving the simulation for modus ponens, the students see that the perceived form of the constants matches the simulated form of the inference rule, indicating that the simulation can be applied. Running the simulation requires first specializing the variables in the simulated premises with the constants from the problem. The simulation then produces " Y " as an inference and specializes it with its value from the first premise, thereby simulating the correct conclusion, "mouse."

As this example illustrates, the same basic processes that simulate natural entities and events can also simulate formal entities and events. It is worth adding, though, that people often construct nonformal simulations to solve formal problems. For example, mathematicians, logicians, and scientists often construct visual simulations to discover and understand formalisms (e.g., Barwise & Etchemendy 1990, 1991; Hadamard 1949; Thagard 1992). Nonacademics similarly use nonformal simulations to process formalisms (e.g., Bassok 1997; Huttenlocher et al. 1994). Whereas proofs ensure the properties of a formalism, perceptual simulations often lead to its discovery in the first place.

4.2. Implications for evolution and development

Amodal symbol systems require a major leap in evolution. Assuming that nonhuman animals do not have amodal symbol systems, humans must have acquired a radically new form of neural hardware to support a radically new form of representation. Of course, this is possible. However, if a more conservative evolutionary path also explains the human conceptual system, parsimony favors it, all other factors being equal.

Not only do perceptual symbol systems offer a more parsimonious account of how intelligence evolved, they also establish continuity with nonhuman animals. On this view, many animals have perceptual symbol systems that allow them to simulate entities and events in their environment. Such simulations could produce useful inferences about what is likely to occur at a given place and time, and about what actions will be effective. Because many animals have attention, working memory, and long-term memory, they

could readily extract elements of perception analytically, integrate them in long-term memory to form simulators, and construct specific simulations in working memory. If so, then the human conceptual ability is continuous with the nonhuman conceptual ability, not discontinuous.

Where human intelligence may diverge is in the use of language to support shared simulations (e.g., Donald 1991; 1993; Tomasello et al. 1993). Evolution may have built upon perceptual symbol systems in nonhuman primates by adding mechanisms in humans for uttering and decoding rapid speech, and for linking speech with conceptual simulations. Rather than evolving a radically new system of representation, evolution may have developed a linguistic system that extended the power of existing perceptual symbol systems. Through language, humans became able to control simulations in the minds of others, including simulations of mental states. As a result, humans became able to coordinate physical and mental events in the service of common goals. Whereas nonhumans primarily construct simulations individually in response to immediate physical and internal environments, humans construct simulations jointly in response to language about nonpresent situations, thereby overcoming the present moment.

Human development may follow a similar path, with ontogeny loosely recapitulating phylogeny (K. Nelson 1996). Similar to nonhumans, infants may develop simulators and map them into their immediate world. During early development, infants focus attention selectively on aspects of experience, integrate them in memory, and construct simulators to represent entities and events (cf. Cohen 1991; Jones & L. Smith 1993; J. Mandler 1992; L. Smith et al. 1992). Long before infants use language, they develop the ability to simulate many aspects of experience. By the time infants are ready for language, they have a tremendous amount of knowledge in place to support its acquisition. As they encounter new words, they attach them to the relevant simulators. New words may sometimes trigger the construction of a new simulator, or a new aspect of an existing one (cf. E. Markman 1989; Tomasello 1992). Much of the time, however, new words may map onto existing simulators and their parts. As linguistic skill develops, children learn to construct simulations productively from other people's utterances, and to construct utterances that convey their internal simulations to others.

Analogous to perceptual symbol systems being continuous across evolution, perceptual symbol systems are continuous across development, with the addition of linguistic control added to achieve social coordination and cultural transmission. Certainly, perceptual symbol systems must change somewhat over both evolution and development. Indeed, the principle of variable embodiment implies such differences (sect. 3.3). To the extent that different animals have different perceptual and bodily systems, they should have different conceptual systems. Similarly, as infants' perceptual and bodily systems develop, their conceptual systems should change accordingly. Nevertheless, the same basic form of conceptual representation remains constant across both evolution and development, and a radically new form is not necessary.

4.3. Implications for neuroscience

Much more is known about how brains implement perception than about how they implement cognition. If

perception underlies cognition, then what we know about perception can be used to understand cognition. Neural accounts of color, form, location, and movement in perception should provide insights into the neural mechanisms that represent this information conceptually. Much research has established that mental imagery produces neural activity in sensory-motor systems, suggesting that common neural mechanisms underlie imagery and perception (e.g., Crammond 1997; Deschaumes-Molinari et al. 1992; Farah 1995; Jeannerod 1994; 1995; Kosslyn 1994; Zatorre et al. 1996). If perceptual processing also underlies cognition, then common sensory-motor mechanisms should be active for all three processes. As described earlier (sect. 2.3), increasing neuroscientific evidence supports this hypothesis, as does increasing behavioral evidence (e.g., Barsalou et al., in press; Solomon & Barsalou 1999a; 1999b; Wu & Barsalou 1999).

Because perception, imagery, and cognition are not identical behaviorally, their neuroanatomical bases should not be identical. Theorists have noted neuroanatomical differences between perception and imagery (e.g., Farah 1988; Kosslyn 1994), and differences also certainly exist between perception and cognition. The argument is not that perception and cognition are identical. It is only that they share representational mechanisms to a considerable extent.

4.4. Implications for artificial intelligence

Modern digital computers are amodal symbol systems. They capture external input using one set of sensing mechanisms (e.g., keyboards, mice) and map it to a different set of representational mechanisms (e.g., binary strings in memory devices). As a result, arbitrary binary strings come to stand for the states of input devices (e.g., 1011 stands for a press of the period key).

Nevertheless, a perceptual symbol system could be implemented in a current computer. To see this, imagine that a computer has a set of peripheral devices that connect it to the world. At a given point in time, the "perceptual state" of the machine is the current state of its peripheral devices. If the machine can focus selectively on a small subset of a perceptual state, associate the subset's components in memory to form a perceptual symbol, integrate this memory with related memories, and later reproduce a superimposition of them in the peripheral device to perform conceptual processing, it is a perceptual symbol system. The machine stores and simulates perceptual symbols in its perceptual systems to perform basic computational operations – it does not process transduced symbols in a separate central processor.

In such a system, a perceptual symbol can later become active during the processing of new input and be fused with the input to form a type-token proposition (in a peripheral device). If different perceptual symbols become fused with the same perceptual input, different construals result. The system could also implement productivity by combining the activation of peripheral states in a top-down manner. Most simply, two symbols could be activated simultaneously and superimposed to form a more complex state never encountered.

Variable embodiment has implications for such implementations. Because peripheral devices in computers differ considerably from human sensory-motor systems, implementations of perceptual symbol systems in computers

should have a distinctly nonhuman character. Contrary to the claims of functionalism, computers should not be capable of implementing a human conceptual system, because they do not have the requisite sensory-motor systems for representing human concepts. Although it is intriguing to consider how a perceptual symbol system might be implemented technologically, it is probably naive to believe that such a system would correspond closely to human intelligence. Such correspondence awaits the development of artifacts that are much more biological in nature.

5. Conclusion

Once forgotten, old ideas seem new. Perhaps the perceptual approach to cognition has not been altogether forgotten, but it has appeared so strange and improbable in modern cognitive science that it has been relegated to the periphery. Rather than resurrecting older perceptual theories and comparing them to amodal theories, reinventing perceptual theories in the contexts of cognitive science and neuroscience may be more productive. Allowing these contexts to inspire a perceptual theory of cognition may lead to a competitive and perhaps superior theory.

Clearly, every aspect of the theory developed here must be refined, including the process of schematic symbol formation, the construct of a simulator, the productive use of language to construct simulations, the fusing of simulations with perceived individuals to produce propositions, the ability to represent abstract concepts, and so forth. Ideally these theoretical analyses should be grounded in neural mechanisms, and ideally they should be formalized computationally. Furthermore, a strong empirical case needs to be established for myriad aspects of the theory. Again, the goal here has only been to demonstrate that it is possible to ground a fully functional conceptual system in sensory-motor mechanisms, thereby giving this classic theory new life in the modern context.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation Grants SBR-9421326 and SBR 9796200. I am indebted to Karen Olseth Solomon, Ling-Ling Wu, Wenchu Yeh, and Barbara Luka for their contributions to the larger project in which this paper originates, and to Jesse Prinz for extensive guidance in reading the philosophical literature. For helpful comments on earlier drafts, I am grateful to Katy Boerner, Eric Dietrich, Shimon Edelman, Arthur Glenberg, Robert Goldstone, Annette Herskovits, Alan Malter, Jean Mandler, Robert McCauley, Douglas Medin, Gregory Murphy, Lance Rips, and Steven Sloman, none of whom necessarily endorse the views expressed. I am also grateful to Howard Nusbaum for discussion on cognitive penetration, and to Gay Snodgrass for permission to use drawings from the Snodgrass and Vanderwart (1980) norms. Finally, I am grateful to Stevan Harnad for supporting this article, for extensive editorial advice of great value, and for coining the term “simulator” as a replacement for “simulation competence” in previous articles.

NOTES

1. This is *not* a claim about correspondence between perceptual symbols and the physical world. Although the structure of perceptual symbols may correspond to the physical world in some cases, it may not in others. For example, philosophers have often argued that a correspondence exists for primary qualities, such as shape, but not for secondary qualities, such as color (e.g., K. Lehrer 1989). Neuroscientists have similarly noted topographic

correspondences between neuroanatomical structure and physical structure (e.g., Tootel et al. 1982; Van Essen 1985).

2. Throughout this paper, double quotes signify words, and italics signify conceptual representations, both modal and amodal. Thus, “chair” signifies the word chair, whereas *chair* signifies the corresponding concept.

3. Some researchers argue that visual agnosia and optic aphasia support a distinction between perception and conception. Because these disorders are characterized by normal perceptual abilities but damaged conceptual and semantic abilities, they suggest that perceptual and conceptual abilities reside in distinct neural systems. Detailed study of these disorders, however, suggests caution in drawing such conclusions (e.g., Hillis & Caramazza 1995). When careful behavioral assessments are performed, correlated deficits between perception and conception may actually be present. Also, bottom-up control of sensory-motor areas may remain after top-down control is lost (sect. 2.4.6). Rather than providing evidence for two different representational systems, these disorders may provide evidence for two different ways of activating a common representational system.

4. Various causal accounts of symbol grounding have been proposed in the philosophical literature, but these typically apply only to a small fragment of concepts and fail to provide a comprehensive account of how symbols in general are grounded. Furthermore, empirical evidence for these theories is typically lacking.

5. Specifying the features computed in the sensory-motor cortices constitutes an undeveloped aspect of the theory. To a considerable extent, this problem belongs to a theory of perception (but see Schyns et al. 1998). Nevertheless, whatever features turn out to be important for perception should also be at least somewhat important for cognition. Specifying how associative areas store patterns of features in sensory-motor areas constitutes another undeveloped aspect of the theory, although the connectionist literature suggests many possibilities.

6. When conscious states do occur, they do not necessarily exhibit one-to-one mappings with the neural states that produced them (Pessoa et al. 1998).

7. Specifying how the cognitive system knows where to focus attention during the symbol formation process constitutes an undeveloped component of the theory. Barsalou (1993) and Logan (1995) suggest several possibilities, but this is a central issue that remains far from resolved in any theory.

8. Jesse Prinz suggested this account of *triangle*. Note that the qualitative neurons in this account constitute a modal representation, not an amodal one. As defined earlier (sect. 1.1), *modal* refers to the fact that the same neurons represent triangles perceptually *and* conceptually. Thus, the qualitative neurons that represent *triangle* are not arbitrarily linked to the neural states that arise while perceiving triangles. Instead, the neurons that represent *triangle* conceptually are a subset of those that are active when triangles are processed perceptually.

9. Recent research suggests that object-oriented reference frames may not be essential to categorization (e.g., S. Edelman 1998; Tarr & Pinker 1989; but see Biederman & Gerhardstein 1993). Regardless, the proposal here is that object-oriented reference frames organize *knowledge* about a type of entity. Although such reference frames may not always become active during familiar categorizations, they almost certainly exist in categorical knowledge, as suggested by people’s robust ability to construct three-dimensional images and perform transformations on them (see Finke, 1989, for a review).

10. Section 2.5 provides preliminary accounts of the frame formation and simulation processes. Nevertheless, many crucial aspects of these processes remain undeveloped, including (1) the integration of memories into frames, (2) the retrieval of information from frames to construct simulations, (3) the integration and transformation of information in simulations, and (4) the development of abstractions.

11. Why cognitive systems divide the world into some categories but not others remains largely unresolved, as in other theo-

ries of knowledge. One exception is so-called basic level categories. Because it is easiest to superimpose perceptual representations of entities that have the same configurations of parts (Biederman 1987; Rosch et al. 1976; B. Tversky & Hemenway 1985), perceptual symbol systems predict that basic level categories should be particularly easy to learn and process (Fig. 3). Also, perceptual symbol systems naturally explain how ad hoc categories arise in cognition (sect. 3.4.4).

12. The criteria for a simulation providing a satisfactory fit to a perceived entity remain unresolved in this theory.

13. As in most other theories of knowledge, how the right inferences are drawn at the right time remains an unresolved issue. Situational cues, however, are likely to be central, with particular inferences activated in situations to which they are relevant. Standard associative mechanisms that implement contextual effects in connectionism may be important, as may the use of affordances described next.

14. Several caveats must be made about the theoretical depictions in Figure 3. For simplicity, this figure only depicts perceptual information in two dimensions, although its actual representation is assumed to be three-dimensional. The object-centered reference frame is also not represented, although some sort of scheme for representing spatial information is assumed. Finally, the depictions in Figure 3 should not be viewed as implying that information in a frame is represented pictorially or as conscious mental images. Again, the actual representations in a frame are configurations of neurons in perceptual systems that become active on representing the physical information conveyed in these drawings. It is essential to remember that these drawings are used *only* for ease of illustration, and that they stand for unconscious neural representations.

15. Framing and contextualized meaning are closely related to the philosophical debate on meaning holism, molecularism, and atomism (e.g., Cummins 1996; Fodor & LePore 1992; Quine 1953). The position taken here is closest to molecularism, which is the idea that concepts often exhibit local dependencies, rather than being globally interrelated (holism), or completely independent (atomism). Specifically, I assume that dependencies reside mostly within simulators and less so beyond. A concept represented recursively within a simulator often depends on other concepts in the same simulator but is relatively less affected by concepts in other simulators.

16. Detailed relations between linguistic and conceptual simulators remain undeveloped in this theory. However, cognitive linguists make many suggestions relevant to developing these relations.

17. These thicker boundaries, too, are theoretical notations that *do not* exist literally in cognitive representations. Instead, they *stand for* the cognitive operation of attention on perceptual representations of space. For a similar view of the role that attention plays in conceptualization and semantics, see Langacker (1986).

18. Obviously, much remains to be learned about when perceptual symbols can and cannot combine productively, and about when emergent features arise and why. Although these issues remain unresolved in all theories of knowledge, constraints from embodiment and affordances are likely to play important roles.

19. Although the focus here is on propositions that bind types to individuals in the world, it is important to note that other kinds of propositions exist as well. For example, types can construe simulators, simulations, operations on simulations, and so forth. The

mechanisms described later for representing abstract concepts in perceptual symbol systems have much potential for handling these other cases.

20. As noted earlier, specifying how the fit between conceptual and perceptual representations is computed constitutes an undeveloped aspect of the theory.

21. Specifying how perceptual symbol systems use external factors and internal content to establish reference constitutes a relatively undeveloped aspect of the theory, as in other theories of symbolic function.

22. Contrary to standard Fregean analysis, senses here function as psychological representations rather than as ideal descriptions that exist independently of human observers.

23. As discussed earlier for intentionality (sect. 3.2.8), resemblance is neither necessary nor sufficient to establish the reference of a perceptual symbol, with external factors playing critical roles. Under many conditions, however, transformations of perceptual symbols do correspond to analogous transformations in their referents, as described here. The conditions under which analogical reference holds remain to be determined.

24. If simulators were sufficiently schematic, such differences might not occur, given that the differences between referents could be completely filtered out. To the extent that the symbol formation process fails to filter out all idiosyncratic details, however, simulators will differ. Widespread exemplar effects suggests that idiosyncratic details often do survive the schematization process (e.g., Brooks 1978; Heit & Barsalou 1996; Medin & Schaffer 1978; Nosofsky 1984).

25. The purpose of this analysis is simply to provide a sense of how perceptual symbol systems could represent abstract concepts. Obviously, much greater detail is needed for a full account. In particular, the external and internal events that frame the simulation must be specified adequately, as must the comparison process, and the assessment of fit. Similar detail is necessary for all the other abstract concepts to follow.

26. If all simulators in long-term memory compete unconsciously to categorize an individual, then at this level a very large number of false propositions are represented, assuming that most simulators do not become bound to the individual successfully. However, one might want to limit the false propositions in cognition only to those considered consciously. In other words, a false proposition becomes represented when a perceiver explicitly attempts to map a simulator into an individual and fails.

27. The alternating simulation in Figure 8 implicitly represents *exclusive or*, given that the horse and cow are never simulated simultaneously, in the middle region. If a third specialization containing both the horse and the cow alternated with the individual specializations, this three-way alternation would implicitly represent *inclusive or*.

28. Simulators and simulations that construe the individuals in Figure 9B as *chair*, *table*, *iron*, and so forth are omitted for simplicity.

29. Kosslyn (1987; 1994) suggests that the right hemisphere represents detailed metric information, whereas the left hemisphere represents qualitative categorical information. To the extent that this formulation is correct, it suggests the following conjecture: the right hemisphere represents individuals in perception and comprehension, and the left hemisphere represents the simulators that construe them. Thus, a type-token proposition involves a mapping from a simulator in the left hemisphere to an individual in the right hemisphere.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Modality and abstract concepts

Fred Adams^a and Kenneth Campbell^b

^aDepartment of Philosophy and ^bDepartment of Psychology, University of Delaware, Newark, DE 19716. {fa; kencamp}@udel.edu
www.udel.edu/Philosophy/famain.html
www.udel.edu/psych/fingerle/Kencamp.htm

Abstract: Our concerns fall into three areas: (1) Barsalou fails to make clear what simulators are (vs. what they do); (2) activation of perceptual areas of the brain during thought does not distinguish between the activation's being constitutive of concepts or a mere causal consequence (Barsalou needs the former); and (3) Barsalou's attempt to explain how modal symbols handle abstraction fails.

Accounts of concepts, cognition, and parsimony pique our interests. Were the brain to use but one kind of symbol (perceptual) for all cognitive tasks, cognition would be economical. Barsalou claims that perceptual symbols (modal symbols) can do whatever nonperceptual symbols (amodal symbols) were posited to do (Barsalou, sect. 1.2.2). If true, cognition would be economically and optimally efficient. However, we are not yet convinced that modal symbols can do everything amodal symbols can.

In what follows, we take up why we believe that Barsalou's account is not yet adequate. We believe that his description of simulators is significantly incomplete. Furthermore, data recording activation of perceptual areas during thinking are consistent with these areas being activated *because* of thinking, not as *part* of thinking. Finally, we maintain that Barsalou's account of simulators is far from showing how abstraction is possible.

Simulators must perform the function of concepts. But what are simulators? Barsalou says what they do, not what they are. He borrows much of what simulators need to do from considering features of cognition that traditionally led to the notion of a concept. Concepts are for categorization, recognition (of token as of type), productivity and systematicity of thought (Aizawa 1997; Fodor & Pylyshyn 1988), and for abstraction and compositional features of thought, generally (Barsalou, sect. 1.5). (Call these cognitive roles R1 . . . Rn). But what is a simulator and how are simulators able to do R1 . . . Rn? This is not clear. Barsalou often helps himself to the very activities traditional concepts are thought to perform without telling us what simulators are or how they are able to do this. Now, to be sure, this target article is not supposed to contain all of the answers. But we would like to see some potential for a future reductive style explanation. It would be music to hear that simulators have properties P1 . . . Pn, and that having P1 . . . Pn explains cognitive roles R1 . . . Rn (the traditional roles for concepts). Instead, we see Barsalou helping himself to R1 . . . Rn without telling us what simulators are or how they produce these properties.

A second significant feature of Barsalou's account is his appeal to neural evidence showing that when people are thinking (not seeing), perceptual or motor regions of the brain become active due to massive cerebral interconnectivity (sects. 2.3, 2.4.7, 4.3). This suggests to Barsalou that perceptual or motor symbols in the brain are constitutive of the concepts employed in thinking. So Fred might think "Ken is in his office" and perceptual regions might be activated that would have previously been used to process the sight of Ken in his office. We do not dispute the evidence that thinking activates perceptual or motor regions. But

Barsalou needs more (sect. 2.1). He needs to demonstrate that these neural connections are *constitutive* of the relevant concepts (of Ken, and of being in one's office) in order to show that concepts do not shed their perceptual roots or ties. Otherwise it is open to the claim that amodal concepts may, when activated, *cause* perceptual or motor symbols to be activated without those symbols being *constitutive* of the related concepts. Fred's concept of Ken is *of Ken*. Thinking of Ken may activate Fred's perceptual symbol for mustache, because Ken has one. But having a mustache is not part of what it is *to be Ken* (he could "lose the mustache"). When we think of *pencils*, motor areas of the brain may be activated for manipulating the hand. But although people can learn to write with their feet (due to loss of hand function), these motor activations cannot be constitutive of our concept of *pencil*. They may just be empirical associations (Adams & Aizawa 1994).

Finally, we turn to Barsalou's account of abstraction. To begin, consider Descartes' classical examples of chiliagons or myriagons. How are our concepts of these figures perceptual (modal)? They seem to be totally dependent on nonmodal definitional understandings (thousand-sided closed figure, million-sided closed figure) respectively. We would be interested in hearing Barsalou's account of the modality of these concepts (or of an infinity of parallel lines in a Lobachevskian space). Our personal simulators lack perceptual ties (symbols) for these concepts.

When Barsalou applies his machinery to abstract concepts, we find no more success. For *truth* (sect. 3.4.3), we are supposed to run a simulator of a balloon above a cloud. This will then be compared with a perception of an actual balloon above a cloud. A match or mapping will elicit from an agent "It's true that a balloon is above a cloud," with *truth* grounded in the mapping (Barsalou, sect. 2.4.6). Supposedly, the agent would acquire the concept of truth over disparate sequences of such mappings and with different modalities. We see two rather large problems with this account. First, Barsalou's evidence supports only that the agent has the concept of *matching* not of *truth*. If AI runs a simulator of what a fork is and then looks at a fork and sees a mapping, this would be as good an account of AI's concept of *fork* (or *match*) as of his concept of *truth*. So, we take it that Barsalou's simulator account of *truth* has important lacunae. Second, Barsalou's appeal to the agent's utterance is disturbing. The meaning of the English word "truth" is carrying some of the load in Barsalou's explanation. But he is *not* entitled to help himself to the meaning of the English word to explain the agent's acquisition of the concept of truth. Meanings of concepts come first. The word comes later (Dretske 1981). Also this threatens to blur derived versus underived meaning for concepts (Dretske 1988; Searle 1983). The meaning of "truth" in a natural language derives from the meaning of truth in the concepts of speakers of the language. Barsalou in several places exploits the use of English sentences to express the contents of perceptual concepts. The problem, of course, is that English (or any natural language) may smuggle meaning from amodal concepts into Barsalou's account.

What makes perceptual symbols perceptual?

Murat Aydede

The University of Chicago, Department of Philosophy, Chicago, IL 60637.
m-aydede@uchicago.edu humanities.uchicago.edu/faculty/aydede/

Abstract: Three major attempts by Barsalou to specify what makes a perceptual symbol perceptual fail. One way to give such an account is to use symbols' causal/nomic relation to what they represent, roughly the way contemporary informational psychosemanticists develop their theories, but this fails to draw the distinction Barsalou seems to have in mind.

Barsalou seems to suggest the following answer(s) to the question of what makes perceptual symbols perceptual: perceptual symbols (PSs) are (1) modal, (2) analogical (i.e., in many cases they re-

semble their semantic value; e.g., referent), and/or (3) reside (are implemented) mainly in the sensory-motor areas of the brain.

The first answer is not helpful because the modal/amodal dichotomy is defined in terms of the perceptual/nonperceptual dichotomy. The most one can get from the text is that “[PSs] are represented in the same systems as the perceptual states that produced them” (sect. 1.1, para. 2). This does not help much either, because all we are told is that a PS is *any* symbol in whose causal production a perceptual state essentially figures, and which resides in the same system as the symbols underlying the perceptual state. For this to work, we must have a workable notion of perceptual/nonperceptual systems, and I am not sure we do, independent of a corresponding notion of perceptual/nonperceptual symbols. Hence it looks like defining modal/amodal symbols in terms of the systems in which they reside will not break the circularity. Furthermore, if causal production by a perceptual state were sufficient to make a symbol perceptual, almost all symbols posited by the recent tradition Barsalou opposes would be perceptual: in a sense, all symbols could be causally produced by perceptual states – if “causal production” is understood as “involvement in causal etiology.”

The second answer is not really elaborated in the text but is confined to some short remarks, such as that the “structure of a perceptual symbol corresponds, at least somewhat, to the perceptual state that produced it” (sect. 1.1). Later in the article, Barsalou tells us a bit more but not much; not enough, anyway, to calm the obvious worries such terminology provokes. He says, for instance, that perceptual symbols *resemble* their referents (sect. 3.2.8), and that certain kinds of changes in the *form* of the symbols systematically result in changes in their semantic value, that is, in what they represent (sect. 3.3). He illustrates this sense in his Figure 6. But he never tells us how to understand such language. This is worrisome, because such discourse borders on the absurd when taken literally. Let me elaborate.

Mental representations (symbols) live dual lives. They are realized by physical properties in the brain. As such, they are physical particulars, that is, they can in principle be individuated in terms of brute-physical properties. Indeed, Barsalou repeatedly reminds us that the pictures in his figures are nothing but a convenient notation to refer to certain activation patterns of neurons in the sensory-motor areas of the brain. *These are* the perceptual symbols. We can, then, talk about symbols as physical entities/events/processes realized in the brain. As such they have a form. But symbols also represent: they have a semantic content. Their content or semantic value is what they represent. The duality in question, then, is one of form/content, or as it is sometimes called, syntax/semantics. PSs must live similar lives. Barsalou’s second answer to my initial question seems to be an attempt to mark the modal/amodal distinction in terms of the *form* of the symbols, or the syntactic *vehicles* of semantic content.

Well, then, how is the attempt supposed to go? How are we to translate his language of “resemblance” so that his answer becomes plausible? What is it about the *formal* properties of vehicles (recall: brute-physical properties of neuronal activation patterns) such that they are *not* semantically “arbitrary”? We are not given any clue in the text. Let me suggest an answer on Barsalou’s behalf.

Perhaps the formal property *F* of a certain type of symbol *S* realized in my brain is semantically nonarbitrary in that its tokening is under the causal/nomic control of instantiations of a certain property of a type of physical object that *S* refers to. But if the answer to our question is to be elaborated along anything like this line, we have abandoned the project of specifying modal/amodal distinction in purely *formal* terms, and have instead causally reached out to the world, to the represented. It is not accidental that such a “reaching out” has been at the core of recent philosophical literature on how to naturalize psychosemantics, in attempts to give a naturalistic account of how certain brain states or neuronal events can represent things/events in the world and have semantic content. This is a matter of the relational/semantic lives

of symbols, not their intrinsic/formal lives, from which I started on behalf of Barsalou.

For Dretske (1981) and Fodor (1987; 1990), two leading informational psychosemanticists, a concept, understood as a symbolic representation realized in the brain, is (partly) the concept it is because it stands in a certain systematic nomological/causal relation to its semantic value. If such relations were crucial to making symbols perceptual, then on this account, many informational theorists would automatically turn out to be defenders of PSs – not a desirable consequence.

I do not think that any general attempt to characterize what makes a mental symbol perceptual/modal (as opposed to nonperceptual/modal) can be successfully given in terms of the formal properties of vehicles. This is plausible if we firmly keep in mind what we are talking about when we talk about mental symbols: activation patterns of neurons in sensory-motor areas of the brain with a lot of physiological properties.

This leaves us with Barsalou’s third answer to my commentary’s question. Perhaps this is indeed his main claim: (almost) all human cognition can be and is done by the kinds of symbols (whatever their forms are) residing in the sensory-motor areas of the brain. But if this is Barsalou’s answer, it does not establish any distinction. For one thing, it is open to the defender of the necessity of amodal symbols in cognition to claim that amodal symbols, for all we know, may reside in such areas. For another, if this is *all* that the perceptual/nonperceptual distinction comes to, the target article loses much of its theoretical interest and bite.

Barsalou’s project is to establish a very strong and controversial claim, namely, that all concepts can be exhaustively implemented by mental symbols which are in some important sense exclusively perceptual. Without an independent and noncircular account of what this sense is, of what it is that makes symbols exclusively perceptual, I am not sure how to understand this claim, let alone evaluate its truth. This is noted in the spirit of putting a challenge rather than an insurmountable difficulty, because I know that Barsalou has a lot to say on this (and on other issues I would have touched on had I more space). Barsalou’s target article contains extremely rich and important material, not all of it clearly formulated and well defended. I very much hope that the target article will (at least) create the opportunity to remedy this weakness in the future.

ACKNOWLEDGMENT

Many thanks to Jesse Prinz and Philip Robbins for their helpful comments.

Perceptual symbols: The power and limitations of a theory of dynamic imagery and structured frames

William F. Brewer

Department of Psychology, University of Illinois, Champaign, IL 61820.
w-brewer@uiuc.edu www.psych.uiuc.edu/~brewer/

Abstract: The perceptual symbol approach to knowledge representation combines structured frames and dynamic imagery. The perceptual symbol approach provides a good account of the representation of scientific models, of some types of naive theories held by children and adults, and of certain reconstructive memory phenomena. The ontological status of perceptual symbols is unclear and this form of representation does not succeed in accounting for all forms of human knowledge.

Importance of perceptual symbols. Much current cognitive psychology is based on amodal propositional representations. Barsalou’s target article is a powerful attempt to emphasize the importance of perceptually based forms of representation. In a recent paper Pani (1996) reviews the history of psychological approaches to imagery. He argues that American functionalists (e.g., Angell, Calkins, Hollingworth, James, Kuhlmann) rejected static “picture

in the head” approaches to imagery and proposed that images were symbolic, abstract, active, and schematic. Pani states that during the behaviorist era this tradition was forgotten so the initial work in cognitive psychology (e.g., Bugelski, early Pavio) tended to take a more static approach, but that current work in imagery is evolving into a view much like that of the early American functionalists. I think that Barsalou’s paper can be seen as setting a new high water mark in this tradition!

There are some domains where I think Barsalou’s perceptual symbol approach can have a major impact. A number of philosophers (Black 1962b; Harré 1970; Hesse 1967) have argued that it is models that give a scientific theory the ability to explain phenomena and to extend the theory to cover new phenomena. These philosophers of science often discussed scientific models in terms of images and were attacked by other philosophers who argued that images did not have the needed representational power (cf. Brewer, in press, for a review). I think Barsalou’s perceptual symbols provide almost exactly the right form of psychological representation for scientific models. Another related area where this theory should play an important role is the study of naive theories in children and adults. For example, in our studies of children’s knowledge of observational astronomy (Vosniadou & Brewer 1992; 1994), we found that children held a wide variety of theories about the day/night cycle (e.g., the sun and moon move up and down parallel to each other on opposite sides of the earth). Perceptual symbols seem to be just the right form of representation to account for our data.

Some unpublished data of mine also seems to bear on the perceptual symbol hypothesis. The Brewer and Treysen (1981) study of memory for places is generally considered a classic demonstration of schema operating in memory. However, some years ago, while thinking along lines parallel to some of those in Barsalou’s target article, I did an analysis of the order in which information was recalled in the written protocols from a room memory experiment. Physical proximity of objects to each other in the observed room proved a better predictor of the *order* of written recall than did conceptual relatedness. This suggests that the recalls were being generated from a perceptual symbol form of representation of the room and not from an abstract, amodal, schema representation.

Ontological status of perceptual symbols. In his earlier papers Barsalou (1993; Barsalou et al. 1993) seems to consider perceptual symbols to be consciously experienced entities (for example, in Barsalou et al. 1993 he refers to his view as “the experiential approach”). In the target article he has made a dramatic shift and states that perceptual symbols are “neural states.” These two positions have rather different implications. In current cognitive psychology the conscious symbol position is perhaps the more radical since that approach would be completely at odds with modern information processing psychology, which is based on unconscious mental processes. The experiential position makes the strong prediction that there will be conscious imagery (or other phenomenal properties) in all tasks involving perceptual symbols. The data go against the hypothesis; for example, Brewer and Pani (1996) found that many participants show little imagery in certain types of memory retrieval tasks (e.g., “What is the opposite of falsehood?”).

The neural position adopted by Barsalou is consistent with the standard unconscious mental processes approach. (Given that we do not yet have scientific knowledge about the specific “configurations of neurons” that underlie perceptual symbols, I do not think the reductive claim, per se, has much practical import.) The neural approach, unlike the experiential one, can easily deal with the lack of images in tasks involving perceptual symbols. However, the perceptual symbol approach is potentially much more powerful than most current theories in cognitive psychology, which are based on representations in terms of unconstrained, arbitrary, amodal propositions. As Barsalou points out, perceptual symbols are not arbitrary and have constraints imposed on them by the assumption that they share structure with perceptual objects. Thus,

the perceptual symbol approach can be confirmed by showing that the unconscious representations show the same constraints as conscious perceptual representations. There are some tasks where this logic can be applied, but it is hard to see how to do this with very abstract concepts since one no longer knows how to define the perceptual constraints.

Problems with representational imperialism. Barsalou argues that all cognitive representation is carried out with perceptual symbols. A list of things the theory has no convincing way to represent: (1) abstract constructs such as entropy, democracy, and abstract; (2) (non model) scientific theories such as evolution and quantum mechanics; (3) gist recall in abstract sentences (Brewer 1975); (4) logical words such as “but,” “therefore,” and “because”; (5) language form; (6) the underlying argument structure of his own article.

Conclusions. The perceptual symbol approach combines the structured representations found in frames with a dynamic view of imagery to produce a very powerful new form of representation that gives a good account for some (but not all) domains of human knowledge.

Perceptual symbol systems and emotion

Louis C. Charland

Department of Philosophy and Faculty of Health Sciences, University of Western Ontario, London, Ontario, Canada N6A 3K7.

charland@julian.uwo.ca publish.uwo.ca/~charland

Abstract: In his target article, Barsalou cites current work on emotion theory but does not explore its relevance for this project. The connection is worth pursuing, since there is a plausible case to be made that emotions form a distinct symbolic information processing system of their own. On some views, that system is argued to be perceptual: a direct connection with Barsalou’s perceptual symbol systems theory. Also relevant is the hypothesis that there may be different modular subsystems *within* emotion and the perennial tension between cognitive and perceptual theories of emotion.

In his target article, Barsalou says he is concerned with perception in a sense that goes beyond the traditional five sensory modalities. He claims to be interested in “any aspect of perceived experience, including proprioception and introspection” (sect. 2.3, para. 3). Introspective processing is said to involve representational states, cognitive operations, and emotional states (sect. 2.3.1). Aside from this remark, Barsalou says nothing else about emotion. He cites the work of four well-known emotion theorists; notably, Damasio (1994), Johnson-Laird (1983), LeDoux (1996), and Mandler (1975). But he fails to say anything substantial about the relationship between what they say about *emotion* and his case for perceptual symbol systems. Yet there are worthwhile connections to pursue.

Perhaps the most obvious link between Barsalou’s views and emotion theory lies in the fact that many emotion theorists believe these have an important *representational* dimension. This is a common and plausible sense of what counts as “symbolic” in cognitive science. There are variations on that theme in emotion theory. For example, in his neurobiological theory of emotion, Antonio Damasio refers to “perceptual images” and dispositional representations” (Damasio 1994, pp. 96–113). On his side, Joseph LeDoux says emotional feelings have a “representational” dimension (LeDoux 1996, pp. 296–300). A good philosophical example of this symbolic orientation in emotion theory is Robert Gordon. He proposes an account of the “logic” of emotion according to which propositional content is central to emotion: when you “emote,” you emote that *p*, where *p* is a proposition (Gordon 1987). In fact there is a long history in analytic philosophical work on emotion that can be interpreted along those lines. Sometimes the symbolic element is belief (Lyons 1980). Sometimes it is normative judgment (Solomon 1976). The notion of judgment also

plays an important part in a variety of psychological theories of emotion. The best example here is Richard Lazarus, who identifies the appraisal component of emotion with “evaluative judgement” (Lazarus 1991). Most appraisal theories of emotion in psychology involve some such symbolic element.

There are therefore a variety of emotion theories according to which these have an important symbolic dimension. In many cases, that thesis is combined with a second one, namely, that in their symbolic capacity, emotions constitute a distinct information processing *system*, governed by its own special affective principles and regularities, with its own special affective representational posits (Charland 1995). Emotion theories modelled along these lines come in two varieties. A good account of a “cognitive” version is Lazarus (1991). A good account of a “perceptual” version is Zajonc (1980). Damasio (1994) is an excellent example of a theory of emotion that manages to reconcile the main insights of both cognitive and perceptual approaches (Charland 1997).

Much like vision and language, then, emotion can be argued to form a distinct symbol processing system of its own. Depending on the theory in question, that system can be cognitive or perceptual, or a combination of both. This is one aspect of emotion theory that is relevant to Barsalou’s project. Unfortunately, in his discussion of “representational processes” he fails to mention emotion (sect. 2.4.7).

Emotion is also relevant to Barsalou’s project in another way. It is possible to argue there are specialized *perceptual* symbol systems *within* the overall affective perceptual representational dimensions of emotion. That hypothesis can be interpreted to mean there are various “modular” symbolic subsystems *in* emotion (Charland 1996). It can also be interpreted to mean that individual emotions have modular features of their own (de Sousa 1987). LeDoux argues this in the case of fear (LeDoux 1996). A promising example of a perceptual emotion module is the facial expressive dimension of emotion (Ekman 1992). Another is Jaak Panksepp’s hypothalamic 4 command-circuit hypothesis (Panksepp 1982; 1998). These last two hypotheses raise interesting questions about the nature and limits of cognitive penetrability, another interest of Barsalou’s (sect. 2.4.6).

One final aspect of emotion theory Barsalou may want to consider has to do with the tension between cognitive and perceptual approaches in cognitive science. There is in fact a partial, but not exact, mirror image of that debate in emotion theory. Emotion theories are often classified in those terms and in a manner that pits the two approaches against one another. I have argued elsewhere that this alleged incompatibility between cognitive and perceptual theories of emotion is more pernicious than it is justified, and that a comprehensive theory of emotion can reconcile the main insights of both (Charland 1997). There may be lessons here that can inform the manner in which Barsalou frames the debate between cognitive and perceptual approaches. Of course, all of this requires a more detailed look at emotion theory than is possible here. Nevertheless, I hope I have shown that there are important connections to pursue in this regard.

Sort-of symbols?

Daniel C. Dennett and Christopher D. Viger

Center for Cognitive Studies, Tufts University, Medford, MA 02155.
{ddennett; cviger01}@tufts.edu

Abstract: Barsalou’s elision of the personal and sub-personal levels tends to conceal the fact that he is, at best, providing the “specs” but not yet a model for his hypothesized perceptual symbols.

John Maynard Keynes was once asked if he thought in words or pictures. “I think in thoughts,” the great man is reported to have replied. Fair enough, but now what? What kind of things are thoughts, and how do you make ‘em out of brainstuff? Keynes’s

answer nicely alerts us to the dangers of oversimplification and false dichotomy, but is otherwise not much help. Similarly, Barsalou’s alternative answer: “we think in perceptual symbols,” is less informative than it might at first appear. There is something compelling about Barsalou’s proposal that cognitive processes be described in terms of simulators (and simulations) involving modal as opposed to amodal formats or systems of representation. Restoring to serious attention the idea that you don’t need a separate (amodal) symbol system to support cognitive functions is a worthwhile project. Moreover Barsalou has interesting suggestions about features that such a perceptuo-motor system ought to have if the brain, one way or another, is to do the work that needs to be done (“the ability to represent types and tokens, to produce categorical inferences, to combine symbols productively, to represent propositions, and to represent abstract concepts” (sect. 1.2.1), but just stipulating that this is possibly what happens in the brain does not begin to address the hard questions.

What does Barsalou mean by “symbol”? He uses the familiar word “symbol” but then subtracts some of its familiar connotations. This is, in itself, a good and familiar strategy (cf. Kosslyn’s (1980) carefully hedged use of “image,” or for that matter, Fodor’s (1975) carefully hedged use of “sentence”). Once Barsalou’s subtraction is done, however, what remains? It’s hard to say. If ever a theory cried out for a computational model, it is here. He says: “perceptual symbol systems attempt to characterize the *mechanisms* that underlie the human conceptual system” (our emphasis; last para., sect. 2.4.6), but here he simply conflates personal and subpersonal cognitive psychology in equating mechanisms with representations; “an important family of basic cognitive processes appears to utilize a single mechanism, namely, sensory-motor representations” (last para. of sect. 2.4.7). These representations, by being *presupposed* to have the very content of our intuitive mental types, must be implicated in a most impressively competent larger structure about which Barsalou is largely silent. As “specs” for a cognitive system there is much to heed here, but the fact that this is *only* specs is easily overlooked. Moreover, if Barsalou’s perceptual symbols are not (personal level) *thoughts* but sub-personal items of machinery, then the content they might be said to have must be a sort of sub-personal content, on pain of reinstating vicious homuncular fallacies.

Are *sort-of* symbols an advance on sort-of sentences and sort-of pictures? How? By not being amodal, one gathers, but also by being only somewhat, or selectively, modal:

First, diagrams such as the balloon in Figure 4A should *not* be viewed as literally representing pictures or conscious images. Instead, these theoretical illustrations *stand for* configurations of neurons that become active in representing the physical information conveyed in these drawings. (para. 3 of sect. 3.1)

Unless we are missing something, this is an assertion of specs without offering a clue about realization. As such, Barsalou’s proposals do not substantially improve on “pure” phenomenology – leaving all the hard work of implementation to somebody else. Consider:

A simulator becomes increasingly active if (1) its frame contains an existing simulation of the individual, or if (2) it can produce a novel simulation that provides a good fit. (para. 2 of sect. 3.2.1)

Fine; now let’s see a model that exhibits these interesting properties. We want to stress, finally, that we think Barsalou offers some very promising sketchy ideas about how the new embodied cognition approach might begin to address the “classical” problems of propositions and concepts. In particular, he found some novel ways of exposing the tension between a neural structure’s carrying specific information about the environment and its playing the sorts of functional roles that symbols play in a representational system. Resolving that tension in a working model, however, remains a job for another day.

On the virtues of going all the way

Shimon Edelman^a and Elise M. Breen^b

^a*School of Cognitive and Computing Sciences, University of Sussex at Brighton, Falmer BN1 9QH, England;* ^b*Laboratory for Cognitive Brain Mapping, Brain Research Institute, Institute of Physical and Chemical Research (RIKEN), Wako, Saitama 351-01, Japan.*
 shimone@cogs.susx.ac.uk www.cogs.susx.ac.uk/users/shimone

Abstract: Representational systems need to use symbols as internal stand-ins for distal quantities and events. Barsalou's ideas go a long way towards making the symbol system theory of representation more appealing, by delegating one critical part of the representational burden – dealing with the constituents of compound structures – to image-like entities. The target article, however, leaves the other critical component of any symbol system theory – the compositional ability to bind the constituents together – underspecified. We point out that the binding problem can be alleviated if a perceptual symbol system is made to rely on image-like entities not only for grounding the constituent symbols, but also for composing these into structures.

Supposing the symbol system postulated by Barsalou is perceptual through and through – what then? The target article outlines an intriguing and exciting theory of cognition in which (1) well-specified, event- or object-linked percepts assume the role traditionally allotted to abstract and arbitrary symbols, and (2) perceptual simulation is substituted for processes traditionally believed to require symbol manipulation, such as deductive reasoning. We take a more extreme stance on the role of perception (in particular, vision) in shaping cognition, and propose, in addition to Barsalou's postulates, that (3) *spatial* frames, endowed with a perceptual structure not unlike that of the retinotopic space, pervade all sensory modalities and are used to support compositionality.

In the target article too, the concept of a frame is invoked as a main explanatory tool in the discussion of compositionality. The reader is even encouraged to think of a frame as a structure with slots where pointers to things and events can be inserted. This, however, turns out to be merely a convenient way to visualize an entity borrowed from artificial intelligence: a formal expression in several variables, each of which needs to be bound to things or events. An analogy between this use of frames and the second labor of Heracles suggests itself: opting for perceptual symbols without offering a perceptual solution to the binding problem is like chopping off the Hydra's heads without staunching the stumps.

The good news is that there *is* a perceptually grounded alternative to abstract frames: spatial (e.g., retinotopic) frames. The origins of this idea can be traced to a number of sources. In vision, it is reminiscent of O'Regan's (1992) call to consider the visual world (which necessarily possesses an apparently two-dimensional spatial structure) as a kind of external memory. In language, a model of sentence processing based on spatial data structures (two-dimensional activation maps) has been proposed a few years ago (Miikkulainen 1993). In a review of the latter work, one of us pointed out that the recourse to a spatial substrate in the processing of temporal structure may lead to a welcome unification of theories of visual and linguistic representation (Edelman 1994).

From the computational standpoint, such unification could be based on two related principles. The first of these is *grounding the symbols* (Harnad 1990) in the external reality; this can be done by imparting to the symbols some structure that would both help to disambiguate their referents and help manipulate the symbols to simulate the manipulation of the referent objects. This principle is already incorporated into Barsalou's theory (cf. his Fig. 6). The second principle, which is seen to be a generalization of the first one, is *grounding the structures* built of symbols.

In the case of vision, structures (that is, scene descriptors) can be naturally grounded in their distal counterparts (scenes) simply by representing the scene internally in a spatial data structure (as envisaged by O'Regan). This can be done by "spreading" the perceptual symbols throughout the visual field, so that in the repre-

sentation (as in the world it reflects) each thing is placed literally where it belongs. To keep down the hardware costs, the system may use channel coding (Snippe & Koenderink 1992; i.e., represent the event "object A at location L" by a superposition of a few events of the form "object A at location L_i ").

In the case of language, structures do not seem to have anything like a natural grounding in any kind of spatial structure (not counting the parse trees that linguists of a certain persuasion like to draw on two-dimensional surfaces). We conjecture, however, that such a grounding is conceivable, and can be used both for representing and manipulating semantic spaces, and for holding syntactic structures in memory (which needs then to be merely a replica, or perhaps a shared part, of the visual scene memory). To support this conjecture, one may look for a "grammar" of spatial relations that would mirror all the requisite theoretical constructs invented by linguists for their purposes (the "syntactic" approach to vision, popular for a brief time in the 1980s, may have remained barren because it aimed to explain vision in terms of language, and not vice versa). Alternatively, it may be preferable to aim at demonstrating performance based on our idea (e.g., by implementing a version of Barsalou's system in which spatial buffers would play the role of the frames), rather than to argue futilely about theories of competence.

In neurobiology, perhaps the best piece of evidence for a perceptual symbol system of the sort proposed by Barsalou is provided by phantom limb phenomena (Ramachandran & Hirstein 1998). These can set in very rapidly (Borsook et al. 1998), are known to occur in the congenitally limb deficient or in early childhood amputees (Melzack et al. 1997), and may even be induced in normal subjects (Ramachandran & Hirstein 1998). In a beautiful experiment, Ramachandran and Rogers-Ramachandran (1996) superimposed a mirror image of the intact arms of amputees onto the space occupied by a phantom arm, and found that movements of the mirrored intact hand produced corresponding kinesthetic sensations in the phantom hand, even in a subject who had not experienced feelings of movement in his phantom hand for some ten years prior to testing. Likewise, touching of the intact mirrored hand produced corresponding, well-localized touch sensations in the phantom hand. These findings support the idea of the world – somatosensory, as well as visual – serving as an external memory (O'Regan 1992), and suggest a stronger relationship between visuospatial and tactile/proprioceptive representations of "body space" for normal subjects. It is interesting that these representations may be linked, in turn, to the mental lexicon: electromyogram (EMG) responses to words with semantic content relating to pain were found to be significantly different in the stumps of amputees with chronic phantom limb pain, compared to the EMG in the intact contralateral limb in the same subjects (Larbig et al. 1996).

Finally, the phantom phenomenon is not limited to the percept of "body space" but may also be demonstrated in other modalities, notably, in the auditory system (Muhlneckel et al. 1998). All this suggests that perceptual symbols, along with a spatial frame provided by the experience of the external world, may (1) solve the symbol grounding problem and (2) circumvent the binding problem – two apparently not immortal heads of the Hydra that besets symbolic theories of knowledge representation. [See also Edelman: "Representation is Representation of Similarities." *BBS* 21(4) 1998.]

Creativity, simulation, and conceptualization

Gilles Fauconnier

Department of Cognitive Science, University of California, San Diego,
La Jolla, CA 92093. gfaucnner@ucsd.edu

Abstract: Understanding the role of simulation in conceptualization has become a priority for cognitive science. Barsalou makes a valuable contribution in that direction. The present commentary points to theoretical issues that need to be refined and elaborated in order to account for key aspects of meaning construction, such as negation, counterfactuals, quantification or analogy. Backstage cognition, with its elaborate bindings, blendings, and mappings, is more complex than Barsalou's discussion might suggest. Language does not directly carry meaning, but rather serves, along with countless other situational elements, as a powerful instrument for prompting its construction.

When a child plays with sugar cubes and matchboxes, pretending that they are cars and buses, driving them all over the kitchen table, and emitting loud “brrr” noises, there is simulation going on in the everyday sense of the word. And there is also simulation going on in the neurobiological and conceptual sense evoked by Barsalou in his target article. If the child's activity makes sense and provides enjoyment, it is because the brain is activating and running some of the dynamic conceptual schemas linked to cars and buses. The projection of such a dynamic frame overlooks and overrides most of the “essential” features commonly associated with vehicles – size, appearance, roads, internal mechanics, and so on. And language is applied effortlessly to the situation: “Look, Mummy, the car hit the bus. The engine's broken.” The activity is in no way exotic. Very young children master it, and use it to develop their conceptual and linguistic systems. The children are in no way deluded. They can maintain simultaneous representations of objects as sugar cubes and as cars. They can mix frames and say things like “Hey, don't put the car in your coffee!” to the adult picking up the cube.

The human capacity for extracting, projecting, and combining dynamic schemas is anything but trivial. It lies at the heart of conceptual creativity. Yet, scientific accounts of meaning seldom do it justice. Focusing only on necessary and sufficient conditions, or on prototypes, or on statistical inference from reality, will quickly rule out our sugar cube adventure. And by doing so, it will rule out most of the power of language and thought. Words like *car* go far beyond picking out categories of objects. They prompt us to construct mappings and run dynamic schemas that can be wildly different depending on circumstances. Reflect if you will on the contribution of the conceptual world of automobiles to the understanding of expressions like *put a tiger in your tank*, *Maytag is the Cadillac of washing machines*, *If cars were men, you'd like your daughter to marry this one* [a Volvo ad], *If cars had wings, we wouldn't need bridges*.

In stressing the crucial role of simulation in conceptualization, and in proposing the general scheme of perceptual symbols to approach the issue in a neurobiologically plausible way, Barsalou does the cognitive science community a great service, offering to liberate it from the chains of amodal computation. In attempting to cover an extremely wide range of issues, there are many complications that Barsalou does not address. Here are a few of them.

Note first that even the simple example above of children at play remains a problem for a literal application of Barsalou's perceptual symbols theory. The mapping of his Figure 3 will fail for the sugar cube and so will many other aspects of perceptual/conceptual mappings commonly applicable to driving. The key here is selective projection. Creative conceptual projection operates on the basis of partial mapping selecting relevant aspects of a complex conceptual domain (Fauconnier & Turner 1998; Tomasello 1998). This suggests that the requirement of a common simulator (sect. 2.4.5) is too strong. More generally, the presupposition in Barsalou's footnote 2, that for a word there is a corresponding concept, is dubious. Wittgensteinian musings and findings in cognitive semantics (Lakoff 1987) suggest a connection of words to multiple

schemas/simulators with many kinds of connections between them, not a direct word to concept correspondence.

A second challenge is the theoretical understanding of emergent properties. Barsalou is quite right to point out that we can conceive of running chairs and croquet flamingos through simulation rather than by feature composition, and that this reflects affordances captured through the formation process. But what is this formation process? We do not get it for free from the perceptual symbol hypothesis. Efforts are made in Fauconnier and Turner 1998, Coulson 1997, and Langacker 1991 to characterize some high level operations and optimality principles that could yield such emergent structure. The problem applied to conceptual blends generally, including counterfactuals (Tetlock et al. 1996; Turner & Fauconnier 1998), fictive motion (Talmy 1996), and grammatical constructions (Mandelblit 1997) is a difficult one. Interesting theories of neural computation that achieve some of these are explored by the Neural Theory of Language Project (Narayan 1997; Shastri & Graner 1996).

A third issue is how to unleash the power of quantifiers, connectives, negation, and variables in a perceptual symbol system. I think Barsalou is on the right track when he thinks of elementary negation in terms of failed mappings. But the proposal needs to be fleshed out. One virtue of a scheme like Barsalou's is that it naturally makes some of the conception part of perception because perceptual symbols and dynamic frames are activated along with more basic perception. We don't see an object with features x , y , z in 3D position p and then compute that it's a balloon above a cloud. Rather we see it directly, conceptually and perceptually all at once, as a balloon above a cloud. Negation, conceived of as failed mapping, is different, because there is an infinite number of possible failed mappings for any situation, but we don't want the cognitive system to try to register all these possible failures. Instead, elementary negation is typically used to contrast one situation with another that is very close to it. “The balloon is under the cloud” is a more likely negation of “The balloon is above the cloud” than “A typewriter is on a desk,” even though the latter would lead to massive mapping failure in that situation. So the notion of failed mapping needs to be elaborated and constrained. Furthermore, the simulations must be able to run internally in the absence of actual events, and so the failed mappings should also be between alternative simulations. Quantifiers, generics, and counterfactuals pose obvious similar problems for the extension of Barsalou's scheme. One rewarding way to deal with such problems is to study the more complex mappings of multiple mental spaces that organize backstage cognition (Fauconnier 1997; Fauconnier & Sweetser 1996) and the cognitive semantics of quantifiers (Langacker 1991).

Many other points should be addressed here, but cannot be for lack of space. Let me only mention that metaphorical and analogical mappings (Gibbs 1994; Hofstadter 1995; Indurkha 1992; Lakoff & Johnson 1999) play a greater and more complex role than Barsalou suggests, and also that in counterfactual and other conceptual blends we routinely find simulations that have no possible counterparts in the real world (Fauconnier & Turner 1998; Turner & Fauconnier 1998).

An important general point for cognitive scientists is that language does not directly carry meaning. Rather, it serves as a powerful means of prompting dynamic on-line constructions of meaning that go far beyond anything explicitly provided by the lexical and grammatical forms (Fauconnier 1997). Fortunately for humans and unfortunately for scientists, the lightning speed operation of our phenomenal backstage cognition in a very rich mental, physical, and social world is largely unconscious and not accessible to direct observation. Still, that is what needs to be studied in order to really account for the famed productivity and creativity of language and thought (Tomasello 1998).

Spatial symbol systems and spatial cognition: A computer science perspective on perception-based symbol processing

Christian Freksa, Thomas Barkowsky,
and Alexander Klippel

Department for Informatics and Cognitive Science Program, University
of Hamburg, D-22527 Hamburg, Germany.

{freksa; barkowsky; klippel}@informatik.uni-hamburg.de

Abstract: People often solve spatially presented cognitive problems more easily than their nonspatial counterparts. We explain this phenomenon by characterizing space as an *inter-modality* that provides common structure to different specific perceptual modalities. The usefulness of spatial structure for knowledge processing on different levels of granularity and for interaction between internal and external processes is described. Map representations are discussed as examples in which the usefulness of spatially organized symbols is particularly evident. External representations and processes can enhance internal representations and processes effectively when the same structures and principles can be implicitly assumed.

The role of spatial relations for cognition. Neural representations resulting from perception are often organized in *sensoritopic* representations, that is, according to spatial structure manifested in the perceived configuration. As entities are perceived in spatial relation to one another, a representation that preserves spatial relations is obtained without costly transformations. In linking the external and the internal worlds, perception processes make use of spatial organization principles.

A special feature of structure-preserving representations is that the same type of process can operate on both the represented and the representing structure (Palmer 1978). [See also Palmer: "Color, Consciousness, and the Isomorphism Constraint" *BBS* 22(6) 1999.] Thus, events in the external world can be reproduced internally.

According to Barsalou (sect. 2.4), simulations with perceptual symbols should underlie the same restrictions as their corresponding perception processes, as they share the properties for arrangement and order of occurrence with "real" perceptions. From this perspective, space – similar to time – plays a twofold role for cognitive processes: spatial location may be represented as a perceptual symbol (as a result of perception) and space provides organizing structure for perceptual symbol systems.

As spatial structure is relevant to perception across modality boundaries and as it behaves like a modality with respect to structural constraints, space can be viewed as an *inter-modality* combining the advantages of specificity and generality of amodal and modal representations, respectively. This inter-modality property of space may be essential for multi-modal sensory integration and for cross-modal interaction.

Specificity versus generality of representations. "General purpose computers" have been praised for their capability in implementing arbitrary concepts and processes. This capability has been achieved by breaking up the structures of specific domains and reducing everything to very general primitive principles that apply to all domains. The generality of today's computer systems is achieved at the expense of an alienation of these systems from a real environment.

The computer is a "brain" that is only superficially connected to the external world (cf. Barsalou sect. 4.4). In reasoning about the world and in simulating processes in the world, computers carry out operations that do not structurally correspond to operations in the world. To make computers function properly according to the rules of a given domain, the specific domain structure must be mimicked and described in an abstract way in the general-purpose structure. As the domain-specific structure is not given intrinsically but is simulated through expensive computational processes, certain operations are achieved at much higher computational cost than in a more specialized structure.

Conversely, highly specific modal structures as manifested in

sensors can do little but react to the specific stimuli they are exposed to. To be effective, they have to be strongly adapted to their environment. In particular, they are not able to perform cross-modal tasks, unless they have a specific interface to establish this connection.

Thus, spatial structure appears to provide a good combination of the advantages of abstract general and concrete specific levels of representation. In spatial structures, a certain degree of abstraction is possible. At the same time, important relations relevant to our perception, conceptualization, reasoning, and action are maintained.

Spatial structures provide principal advantages for processing information about spatial environments (Glasgow et al. 1995). This can be attributed to the fact that – unlike other information structures – spatial relationships are meaningful on many different granularity levels.

Distinctions that are made on high-resolution levels disappear on lower resolution levels. Preserving spatial structure, different meaningful conceptualization levels can be achieved simply by "looking" more closely or less closely (Freksa 1997).

This property of spatial structure can also be utilized on more abstract levels: conceptual spaces in which small local changes can be ignored in favor of the similarities on a higher level form conceptual neighborhoods (Freksa 1991) with interesting computational properties (Freksa & Barkowsky 1996).

Maps as external representations for cognitive processes. Maps and map-like representations convey knowledge by depicting spatial entities in a planar external medium. Relevant geographic entities are identified and depicted by more or less abstract symbols that form spatially neighboring pictorial entities on the map's surface (Barkowsky & Freksa 1997).

The usefulness of maps depends strongly on the spatial structure of the representational medium and on human spatio-perceptual recognition and interpretation abilities. The symbols in the map are perceived in relation to one another. The pictorial map space is organized in spatial analogy to the world it represents. Hence, the map becomes an interface between spatio-perceptual and symbolic concepts that are integrated in a single representational structure.

Maps as external knowledge representations preserve spatial relations adapted from the perception of the represented spatial environment. Hence they support the conception and use of spatial information in a natural way. The spatial structure of maps extends the cognitive facilities of the map user to the external representation medium (Scaife & Rogers 1996).

The cognitive information processing principles that are applicable to both internal and external spatial representations can be investigated by studying processes on external spatial representation media. This seems especially valid regarding spatially organized knowledge for which the types of relations relevant to both internal and external conceptions of environments can be studied (Berendt et al. 1998). Moreover, as maps combine both spatio-analogical and abstract symbolic information, the relevance of spatial organization principles for non-spatial information can be investigated.

Extending brain power: External memory and process model. The idea of using the same principles for internal and external knowledge organization is computationally appealing. Common structures support the use of internal processes for using externally represented knowledge. Consequently, external knowledge can serve effectively as an extension of internal memory structures.

Furthermore, internal and external processes that strongly correspond to one another can greatly enhance the communication between the internal and the external worlds. A main reason for this is that only a little information needs to be exchanged to refer to a known corresponding process, as the common structures can act implicitly. In anthropomorphic terms we could say that it is possible to empathize with familiar processes but not with unfamiliar ones.

ACKNOWLEDGMENT

We are grateful to Deutsche Forschungsgemeinschaft (Spatial Cognition Priority Program) for support and to Stephanie Kelter for valuable comments.

Grounded in perceptions yet transformed into amodal symbols

Liane Gabora

Center Leo Apostel, Brussels Free University, 1160 Brussels, Belgium.
lgabora@vub.ac.be www.vub.ac.be/CLEA/liane

Abstract: Amodality is not incompatible with being *originally* derived from sensory experience. The transformation of perceptual symbols into amodal abstractions could take place spontaneously through self-organizing processes such as autocatalysis. The organizational role played by “simulators” happens implicitly in a neural network, and quite possibly, in the brain as well.

The symbol grounding problem is important, and Barsalou’s target article does an excellent job of directing our attention to it. However, to be *grounded* in perceptual experience does not mean to spend eternity underground. Just as a seed transforms into a sprout which emerges from the ground as a plant, an abstract concept can originate through, or be grounded in, perceptual experience, yet turn into something quite different from anything ever directly perceived. To insist that abstractions are just arrangements of perceptual symbols is like insisting that a plant is just seed + water + sunlight.

As Barsalou considers increasingly abstract concepts, he increases the complexity of the perceptual symbol arrangements. To me it seems more parsimonious to say that what was once a constellation of memories of similar experiences has organized itself into an entity whose structure and pattern reside primarily at a level that was not present in the constituents from which it was derived. This does not mean that abstractions cannot retain something of their “perceptual character.”

Abstraction can transform sensory experiences into amodal symbols. Barsalou assumes that viewing symbols as amodal is incompatible with their being derived from sensory experience, because of the “failure to provide a satisfactory account of the transduction process that maps perceptual states into amodal symbols” (sect. 1.2.2, para. 4). He admits that perceptual symbols are analogical, however, and that they are integrated “combinatorially and recursively” (sect. 1.5); and he speaks of “filtering out the specific circumstances” (sect. 2.3.1., para. 1). Aren’t these the kind of processes that could transduce perceptual symbols? Incidentally, it is strange that Barsalou cites Gentner, who works on structure mapping in analogy (e.g., Gentner 1983; Gentner & Markman 1997) to support the notion that “mental models are roughly equivalent to only the surface level” and “tend not to address underlying generative mechanisms that produce a family of related simulations” (sect. 2.4.2).

Barsalou pays surprisingly little attention to the highly relevant work of connectionists. The only rationale he provides is to say that because “the starting weights on the connections . . . are set to small *random* values . . . the relation between a conceptual representation and its perceptual input is arbitrary” (sect. 1.2). This is misleading. Just because there are many paths to an attractor does not mean the attractor is not attracting. Surely there are analogous small random differences in real brains. This quick dismissal of connectionism is unfortunate because it has addressed many of the issues Barsalou addresses, but its more rigorous approach leads to greater clarity.

Consider, for example, what happens when a neural network abstracts a prototype such as the concept “depth,” which is used in senses ranging from “deep blue sea” to “deep-frozen vegetables” to “deeply moving book.” The various context-specific inter-

pretations of the word cancel one another out; thus the concept is, for all intents and purposes, amodal, though acquired through specific instances. There is no reason to believe that this does not happen in brains as well. The organization role Barsalou ascribes to “simulators” emerges implicitly in the dynamics of the neural network. Moreover, since Barsalou claims that “a concept is equivalent to a simulator” (sect. 2.4.3), it is questionable whether the new jargon earns its keep. Why not just stick with the word “concept”?

How would amodal cognition get started? Harder to counter is Barsalou’s critique that an amodal system necessitates “evolving a radically new system of representation” (sect. 4.2). Barsalou repeatedly hints at but does not explicitly claim that the origin of abstract thought presents a sort of chicken-and-egg problem. That is, it is difficult to see how an abstraction could come into existence before discrete perceptual memories have been woven into an interconnected worldview that can guide representational re-description down potentially fruitful paths. Yet it is just as difficult to see how the interconnected worldview could exist prior to the existence of abstractions; one would expect them to be the glue that holds the structure together. In Gabora 1998 and Gabora 1999, I outline a speculative model of how this might happen, drawing on Kauffman’s (1993) theory of how an information-evolving system emerges through the self-organization of an autocatalytic network. Self-organizing processes are rampant in natural systems (Kauffman 1993) and could quite conceivably produce a phase transition that catapults the kind of change in representational strategy that Barsalou rightly claims is necessary.

Did the divide between humans and other animals begin with language? It also seems misleading to cite Donald (1991) [See also: multiple book review of Donald’s *Origins of the Modern Mind* BBS 16(4) 1993.] as supporting the statement “where human intelligence may diverge [from animals] is in the use of language” (sect. 4.2). Donald writes:

Speech provided humans with a rapid, efficient means of constructing and transmitting verbal symbols; but what good would such an ability have done if there was not even the most rudimentary form of representation already in place? There had to be some sort of semantic foundation for speech to have proven useful, and mimetic culture would have provided it.

Thus Donald argues that human intelligence diverged from animals *before* the appearance of language, through the advent of mimetic skill.

Despite these reservations, I found the target article interesting and insightful.

Embodied metaphor in perceptual symbols

Raymond W. Gibbs, Jr. and Eric A. Berg

Department of Psychology, University of California, Santa Cruz, Santa Cruz, CA 95064. gibbs@cats.ucsc.edu

Abstract: We agree with Barsalou’s claim about the importance of perceptual symbols in a theory of abstract concepts. Yet we maintain that the richness of many abstract concepts arises from the metaphorical mapping of recurring patterns of perceptual, embodied experience to provide essential structure to these abstract ideas.

We strongly endorse Barsalou’s argument that a perceptual theory of knowledge can implement a fully functional conceptual system. One of Barsalou’s most important claims is that some abstract concepts are partly represented in terms of perceptual symbols, but metaphor is only used to elaborate on and extend these abstract concepts. For instance, our understanding of the concept “anger” is not entirely explained by viewing it in metaphorical terms as “liquid exploding from a container.” Barsalou claims that this metaphor fails to capture the richness of our concept of “anger.”

Instead, direct perceptual symbols underlie the representation of the abstract concept for “anger,” which metaphor may elaborate upon as a secondary process. Direct perceptual representations for abstract concepts are presumably needed to guide the mapping of concrete domains (e.g., “liquid exploding in a container”) into abstract target domains (e.g., “anger”). Thus, abstract knowledge domains are shaped by perceptual representations, but not fundamentally by metaphor.

We take issue with this conclusion. Metaphorical processes may be essential in the formation and perceptual representation of many abstract concepts. For instance, Barsalou notes that people have direct experience of “anger” in three ways. “Anger” first involves the appraisal of the initiating event, such as that a person’s goals have been thwarted. Second, “anger” gives rise to intense, affective states. Finally, “anger” gives rise to significant behavioral responses, such as seeking revenge or working toward removal of the obstacle. These three components of our experience provide the direct perceptual knowledge underlying the formation of perceptual symbols for the concept of “anger.”

Yet it is almost impossible to see these fundamental aspects of our anger experiences apart from the metaphors we tacitly employ to make sense of these experiences. Thus, the idea of “anger is heated fluid in a container” is one of several prominent metaphorical ideas that we have to structure our anger experiences. Central is our understanding of the conceptual metaphor “anger is heated fluid in a container” is the embodied experience of containment. We have strong kinesthetic experiences of bodily containment ranging from situations in which our bodies are in and out of containers (e.g., bathtubs, beds, rooms, houses) to experiences of our bodies as containers in which substances enter and exit. A big part of bodily containment is the experience of our bodies being filled with liquids including stomach fluids and blood, and sweat that is excreted through the skin. Under stress, we experience the feeling of our bodily fluid becoming heated.

Our immediate experience of anger is recognized and understood in terms of, among other things, our direct embodied experience of being containers with fluids that sometimes get heated when put under pressure from some external stimulus. Note the correspondences between different embodied experiences while feeling angry and the language we use to talk about different aspects of the concept for “anger.” Thus, when the intensity of anger increases, the fluid in the bodily container rises (e.g., “His pent-up anger welled up inside of him”). We also know that intense heat produces steam and creates pressure on the container (e.g., “Bill is getting hot under the collar” and “Jim’s just blowing off steam”). Intense anger produces pressure on the container (e.g., “He was bursting with anger”). Finally, when the pressure of the container becomes too high, the bodily container explodes (e.g., “He blew up at me”). Each of these aspects of our folk understanding of the concept of “anger” is a direct result of the metaphorical mapping of heated fluid in a container onto the abstract domain of “anger.”

Several kinds of evidence from psycholinguistics supports the idea that embodied metaphors underlie people’s understanding of abstract concepts (Gibbs 1992; 1994). These studies showed that people were remarkably consistent in their answers to questions about their understanding of events corresponding to particular bodily experiences that were viewed as motivating specific source domains in conceptual metaphors (e.g., the experience of one’s body as a container filled with fluid). For instance, people responded that the cause of a sealed bodily container exploding its contents out is the internal pressure caused by the increase in the heat of the fluid inside the bodily container, that this explosion is unintentional because containers and fluid have no intentional agency, and that the explosion occurs in a violent manner. These responses were provided without the experimenter mentioning to participants the idea of “anger” and as such simply reflect people’s understanding of their perceptual, embodied experiences.

More interesting, though, is that people’s intuitions about various source domains map onto their conceptualizations of different target domains in very predictable ways. For instance, several

experiments showed that people find idioms about anger to be more appropriate and easier to understand when they are seen in discourse contexts that are consistent with the embodied experiences underlying their concept of “anger.” Thus people find it easy to process the idiomatic phrase “blow your stack” when this was read in a context that accurately described the cause of the person’s anger as being due to internal pressure, where the expression of anger was unintentional and violent (all entailments that are consistent) with the entailments of the source to target domain mappings of heated fluid in a container onto anger). But readers took significantly longer to read “blow your stack” when any of these entailments were contradicted in the preceding story context. This evidence suggests that people’s understanding of various linguistic phrases were related to their embodied, metaphorical conceptualizations of different abstract concepts to which idioms refer.

Another important part of the metaphorical nature of perceptual symbols arises from the correlations people note in their perceptual, embodied experience. For instance, consider the abstract concept of “theory,” which is motivated by several conceptual metaphors, including “theories are buildings” (e.g., “Your theory needs more support” or “Your theory can’t hold up to the weight of the existing data”). At first glance, the abstract concept of a theory does not seem to be perceptually related to the buildings in which we generate, discuss, and dismantle them. But several more primitive metaphors underlie our metaphorical understanding of “theories are buildings,” such as the fact that “persisting is remaining erect” and “structure is physical structure” (Grady 1997). These primitive metaphors arise from our perceptual, embodied experience where we note strong correlations between things, including animate objects, being erect and their persisting, and between entities we view as structures and physical structures. Our understanding of the abstract concept for “theory” – theories are buildings” (now called “compound metaphors”) – emerges, in part, from these recurring patterns of correlated, perceptual experiences. However, our rich understanding of abstract concepts must include the metaphorical mapping of these perceptual correlations as source domains onto the target domain of the abstract concept. Under this view, abstract concepts are, as Barsalou argues, represented in terms of perceptual symbols, yet these symbols are, in many cases, inherently structured by metaphor.

Perceptual symbols in language comprehension

Arthur M. Glenberg

Department of Psychology, University of Wisconsin – Madison, Madison, WI 53706. glenberg@facstaff.wisc.edu

Abstract: Barsalou proposes (sect. 4.1.6) that perceptual symbols play a role in language processing. Data from our laboratory document this role and suggest the sorts of constraints used by simulators during language comprehension.

A conundrum: Barsalou develops a strong case that cognition is not based on rule-like (e.g., syntactic) manipulation of abstract amodal symbols. But words in a sentence are quintessentially amodal symbols operated on by syntactic rules. So, how can we understand language? In section 2.6, Barsalou sketches an answer that is almost identical to one that we have developed (Glenberg & Robertson, in press; Glenberg et al., under review) and called the Indexical Hypothesis.

The Indexical Hypothesis specifies three steps in understanding language in context. First, words and phrases are indexed to objects in the environment or to perceptual symbols from memory. Second, affordances are derived from the objects or perceptual symbols. In most ordinary language situations, these affordances are action-based. That is, the affordances are actions that

Table 1 (Glenberg). *Two example scenarios*

Setting: Marissa forgot to bring her <i>pillow</i> on her camping trip.
Afforded: As a substitute for her <i>pillow</i> , she filled up an old sweater with leaves .
Non-afforded: As a substitute for her <i>pillow</i> , she filled up an old sweater with water .
Related: As a substitute for her <i>pillow</i> , she filled up an old sweater with clothes .
Setting: Mike was freezing while walking up State Street into a brisk wind. He knew that he had to get his face covered pretty soon or he would get frostbite. Unfortunately, he didn't have enough money to buy a scarf.
Afforded: Being clever, he walked into a store and bought a newspaper to cover his <i>face</i> .
Non-afforded: Being clever, he walked into a store and bought a matchbook to cover his <i>face</i> .
Related: Being clever, he walked into a store and bought a ski-mask to cover his <i>face</i> .

Note: Central concepts are in italic and distinguishing concepts are in boldface. From Glenberg et al. *Symbol grounding and meaning*.

matchbooks provide protection from the wind. Second, by focusing on embodied constraints on action, we can better understand Barsalou's claim (sect. 2.4.3) that "the primary goal of human learning is to establish simulators." Being able to simulate is unlikely to have been selected by evolution if simulation did not contribute to survival and reproductive success, both of which require action, not just thought. On the other hand, if simulations result in derivation of new affordances, then the simulations related directly to effective action. Third, action-based constraints suggest a revision of some implications for perceptual symbol systems. For example, if the goal of cognition is to support action in the world, and it is action that grounds symbols (e.g., Newton 1996), then the action-less perceptual symbols Barsalou describes in section 4.4 are insufficient to generate intelligence. That is, although the machine would have perceptual symbols, those symbols would not convey to the machine how it should interact with the world it is experiencing. Such a machine could not act intelligently.

Perception as purposeful inquiry: We elect where to direct each glance, and determine what is encoded within and between glances

Julian Hochberg

Department of Psychology, Columbia University, New York, NY 10027
hochberg@columbia.edu

Abstract: In agreement with Barsalou's point that perceptions are not the records or the products of a recording system, and with a nod to an older system in which perception is an activity of testing what future glances bring, I argue that the behavior of perceptual inquiry necessarily makes choices in what is sampled; in what and how the sample is encoded, and what structure across samples is pursued and tested; and when to conclude the inquiry. Much of this is now being hectically rediscovered, but a comprehensive approach like the one Barsalou proposes should help preserve what progress is made.

Barsalou provides what looks like a viable big picture, a perceptual view of cognition, in which everything can take its place, with no crippling simplifications. I want to expand slightly on his assertion that perception is not a recording system. Indeed, percep-

can be accomplished with the object by a person with the same type of body as the comprehender. The particular affordances derived depend on the goals that need to be accomplished, affective states, and grammatical form. As an example, a wooden chair affords actions such as sitting-on, standing-on (e.g., to change a light bulb), and if both the chair and the person are of the right form, a chair can be lifted to afford defense against snarling lions. Thus, the goals of resting, standing, or defense select some affordances over others. The third step in language understanding is to combine, or mesh (Glenberg 1997) the affordances. Meshing is a process of combining concepts (perceptual simulators) that respects intrinsic constraints on combination. These constraints arise because not all actions can be combined, given the way our bodies work. Thus, whereas the affordances of chairs can be meshed with the actions underlying the goals of resting, light bulb-changing, or defense, an ordinary chair cannot be meshed with goal of propelling oneself across town or serving as a vase. Consequently, we can understand a sentence such as "Art used the chair to defend himself against the snarling lion," but we find it difficult to understand, "Art used the chair to propel himself across town." In general, if the affordances can be meshed to simulate coherent action, then the sentence is understood; if the affordances cannot be meshed, then the sentence is not understood.

To test these ideas, participants were asked to judge the sensibility of afforded and non-afforded sentences in the context of a setting sentence (see Table 1). In the afforded condition, the affordances of named objects can be meshed to accomplish stated goals. Thus, sweaters and leaves can be combined to make a pillow. In the non-afforded condition, the affordances cannot be meshed to accomplish the goal. Hence, we predicted that experimental participants will have difficulty understanding the non-afforded sentences. Note that objects such as sweaters and newspapers are used in novel ways in both sentences. Thus, it is highly unlikely that understanders can base their judgment on specific experiences with the objects. In fact, we demonstrated using latent semantic analysis (LSA; Landauer & Dumais 1997) that in tens of thousands of texts, a central concept of each sentence (e.g., "pillow") appears in contexts that are orthogonal to the contexts in which the distinguishing concepts (e.g., "water" and "leaves") appear.

Participants in the experiment rated the afforded sentences as sensible (average of 4.58 on a 7-point scale from 1: virtual nonsense to 7: completely sensible), whereas they rated the non-afforded sentences as pretty much nonsense (average rating of 1.25). One might argue that the ratings reflect syllogistic problem solving rather than reading and comprehension. To answer this argument, we wrote related sentences (see Table 1) in which (a) the distinguishing concepts (e.g., "ski mask") and the central concepts (e.g., "face") were related according to the LSA text analysis, and (b) the affordances would mesh. Thus, the related sentences described situations we thought would be familiar to our participants, whereas the afforded sentences described novel situations. If people were understanding the afforded sentences through syllogistic problem solving, then the afforded sentences should be read more slowly than the related sentences that do not demand problem solving because of their familiarity. In fact, the times taken to read and understand the afforded and related sentences were statistically indistinguishable, whereas both were read faster than the non-afforded sentences. Because comprehension depended on the ability to derive affordances for objects (e.g., that newspapers can afford protection from the wind), we demonstrated that conceptual representations must include something like perceptual symbols that allow such novel derivations. Because not all concepts can be combined, we demonstrated embodied constraints on conceptual combination and sentence understanding.

Action-based constraints on mesh serve three purposes in a perceptual symbol system. First, they are required to ensure the proper operation of perceptual simulators. Otherwise, people might be led to believe that wet sweaters can be pillows and that

tion is a purposive behavior, depending on the viewer's intentions and expectations.

This view is not new. To Helmholtz (1866/1962), the idea of an object was what one would see after changing viewpoint. To J. S. Mill (1865/1965), the idea that an object is real and not imagined reflects the permanent possibility of obtaining sensations from it. Hebb's 1949 *phase sequence* embodied some of the better aspects of these metatheories in something closer to an actual theory. In these and in other approaches (including Barsalou's most comprehensive system) a perception is not a transcribed recording or a given image: it refers to an active inquiry, its findings, and its consequences.

In none of these approaches is the perception composed of *consciously accessible* components, or sensations; in all of these, the purpose or intention of the perceptual inquiry determines what stimulation is attended, and the inquiry itself is a course of behavior over time. Such behavior itself needs guidance, with constraints on what will be acceptable timing and helpful peripheral salience, constraints that have in effect determined much of our visual culture and technology (i.e., of text and display layout, pictorial composition, movie cutting, etc.). (Human factors should of course pay heed to experimental psychology, but psychologists would do even better to study these aspects of the visual media, a point I have pursued elsewhere at length.)

We can note at least three levels of perceptual inquiry, and none is a record: each depends on the viewer's purpose and resources, although they do so in different ways.

1. Looking behavior is in large part about bringing some specific part of the retinal image, *chosen in advance*, close to the foveal region (approximately the size of a thumbnail at half-arm's length): Attentional shifts to specific parts of the field of view do in fact precede the eye movements which make them visible (Hoffman & Subramaniam 1995; Kowler et al. 1995). To the low resolution of peripheral vision, the junctions of a drawn cube, for example, *afford* notice of where occlusion information can be found; similarly, the gaps in text predict where words begin and end. Unless you choose to fixate these, you have no sensory information about which face of the cube is nearer, or what the word's initial letters are in fact, in cube and text respectively. With a directed perceptual task, or with familiar and normally redundant text, one glance may satisfy the perceptual inquiry. Such perceptual "inference" may furnish the wrong answer, as shown by proof-reading errors and by the acceptance of locally inconsistent (Penrose/Escher) figures.

2. Can we then treat perception as a record, complete or otherwise, if not of the optic array then of what is obtained from the momentary glance with its inhomogeneous retinal resolution? Not so: even within a fixation, both the viewer's purpose (or instructions) and the memory labels or structures available for encoding affect what is encoded into working memory within that brief glance (normally about 0.25 sec – 0.5 sec). Whatever has not been so encoded is effectively lost, although Sperling's landmark *partial report* experiments (Sperling 1960) show that that material was briefly available to the viewer for alternative encoding. Some of what has been attributed to inattention blindness most plausibly reflects this forgetting of what was not intentionally encoded (Hochberg 1970; 1988; Wolfe 1998); that point applies as well to multiglance inquiries, as noted next.

3. Does the perception of extended objects and events, which require more than one glance, amount to a record of the encodings that were made in the component glances? Again, not so (except in forced, slow list-checking). When perceptual inquiry spans multiple glances, normally at 2–4 glances/sec, the separate encodings must themselves be encoded within or tested against some superordinate schematic structure from which they can be retrieved, or they too are lost. Both intention and memory are needed to bridge successive glances at even relatively compact objects (Hochberg 1968), layouts (Irwin 1996), or serial aperture viewing (Hochberg 1968); in listening to superimposed speech (Treisman & Geffen 1967, as reinterpreted by Hochberg 1970);

and in viewing superimposed filmed events (Neisser & Becklen 1975) and the elisions that are common in normal film cutting (Hochberg & Brooks 1996).

Not only are perceptions not recordings, but the levels of perceptual inquiry (where one fixates, what one encodes from that glance, whether one goes on to take more glances, and how one encodes the several glances) offer different routes to selection and organization that are inherent in the nature of perceptual inquiry (Hochberg 1998). With all of this cognitive involvement in the perceptual process, it has always seemed strange to stop at the firewall that has separated theories about perception from those about thinking and reasoning for most of this century. Barsalou's creative and encompassing proposal should help provide a comprehensive context within which to examine perception as a cognitive behavior, and to place and hold what such examination reveals.

Individuals are abstractions

James R. Hurford

Department of Linguistics, University of Edinburgh, Edinburgh, EH8 9LL, Scotland, United Kingdom. jim@ling.ed.ac.uk www.ling.ed.ac.uk/~jim/

Abstract: Barsalou's move to a perceptual basis for cognition is welcome. His scheme contrasts with classical logical schemes in many ways, including its implications for the status of individuals. Barsalou deals mainly with perceived individuals, omitting discussion of cognized individuals. It is argued that the individuality of cognized individuals is an abstraction, which conforms in its manner of formation to other cognitive abstractions which Barsalou discusses, such as truth and disjunction.

In his impressive and welcome target article, Barsalou is at pains not to throw out some cherished symbolic babies with the amodal bathwater. Many of these babies, in their best-known incarnations, are the progeny of classical symbolic logic; they include the notions of proposition, truth, and abstraction. Barsalou sketches how the work that such notions do in classical theories can also be done in his theory, and this involves casting these theoretical terms in a new light. Barsalou does not focus specifically on the notion of an individual, but the idea figures persistently in his discussion, and it is clear that he does not want to throw it out. His theory implies that we should take a radically different view from that of classical modern logics as to what individuals (alias particulars) are.

The classical modern logical view (e.g., Carnap 1958) assumes an objectively given domain consisting of individual objects, typically exemplified as concrete (e.g., the sun, the moon, the person Charles). It usually goes without saying that each object in the domain is distinctive and not identical to any other object. Furthermore, the objects in the domain are taken as the most basic elements of meaning. This is evidenced in the postulation by formal semanticists such as Montague (1973) of *e* as a basic ontological type, the type of individual entities. How such objects are known or perceived by people is not the concern of logicians.

The logical form of a specific elementary proposition uses individual constants as arguments of predicates, and emphasizes the different functions in the logical system of predicates and individual constants. When Barsalou begins to discuss propositions (sect. 3.2), his examples are such as ABOVE(ceiling, floor) and BELOW(floor, ceiling). Clearly, these are not the propositions envisaged by logic, because, without being dazzled by the change from upper to lower case type and interpreting floor and ceiling in a natural way, these latter terms are predicates and not individual constants. I can point to something, and say of it (i.e., predicate of it) that it is a floor. When students in an introductory logic class confuse predicates with individual constants, we penalize them. But in the long run, I believe, such students and Barsalou are right, and classical modern logic has got it wrong.

Barsalou writes that "perceptual symbols need not represent

specific individuals . . . we should be surprised if the cognitive system *ever* contains a complete representation of an individual. . . . we should again be surprised if the cognitive system ever remembers an individual with perfect accuracy” (sect. 2.2.3). This again shows the contrast with the logical approach, in which there can be no question of the completeness or accuracy of an individual constant term in denoting an individual; relative completeness is out of the question because an individual constant term is atomic, and accuracy is given by the logician’s fiat.

At his most careful, Barsalou speaks of “perceived individuals” or “individuals in the perceived world,” and this is where his version of proposition is based. “Binding a simulator with a perceived individual . . . constitutes a type-token mapping. . . . This type-token mapping implicitly constitutes a proposition” (sect. 3.2.1). Thus, in the very act of perceiving something as belonging to some pre-established mental category, I can entertain a proposition, because my attention at the moment of perception is fixed on this one particular thing. The argument of the predicate is given by my attention at the moment. The numerical singularity of the argument comes from the bottom up, because I am attending to just one thing; the predication comes from a match between some perceived property of the thing and top down information about a pre-existing mental category. Thus, a Barsalovian proposition is formed.

The numerical singularity of a perceived object is a product of the observer’s attention at the time of perception. An observer may focus on a pile of rice or on a single grain, or on a pair of boots attended to as a pair, or on two separate individual boots. If I attend to two objects, simultaneously but individually (say, a man and his shaving mirror), my perception of them as two individual objects at the time provides slots for two arguments, and I can find a pre-established 2-place mental relation (say, *use*) to classify the perceived event. (The limit on how many individuals one can attend to at once may be similar to the limits of subitization in young children – up to about four items.)

A classic formal semantic model (e.g., Cann 1993) might contain a set of individual entities, each satisfying the predicate *ant*; the denotation of the predicate *ant* is the set. Each individual *ant* is the denotation of some individual constant (e.g., *a*₁, *a*₂, *a*₃, . . .) in the logical language. Thus, *a*₁ and *a*₂ are distinguished by the fiat according to which the logician constructs his model. The proposition *ant* (*a*₁) is a different proposition from *ant* (*a*₂); and no further distinction – for example by a predicate applying to one *ant* but not to another – is necessary in the logical system. The representation of the proposition in the logical language tells us which *ant* the proposition is about.

By contrast, in terms of perception, I can attend to certain properties of an object and judge from these properties that it is an *ant*, but the perceived properties cannot tell me which *ant* it is. I know that the world contains more than one *ant*, because I have sometimes seen many together, but all I know is that I have seen some *ant*. Barsalou’s account of the storing of a basic proposition in long-term memory describes it as involving a “type-token fusion” (sect. 4.1.5). This fusion is not further described, but it must in fact result in the loss of identity of the token, the “whichness” of the originally perceived object.

Barsalou is usually careful to prefix “perceived” onto “individual.” We can see how an individual can be perceived, but can an individual be cognized, and if so, how? If individuals lose their “whichness” during the process of storing a type-token fusion in long-term memory, how can the perceptual symbols (the simulators) in my mind for an *ant* and for my mother differ in a way that echoes the classical difference between a set and an individual (or between a property and an individual concept)?

A cognized individual (as opposed to a perceived individual) can be constructed by the process involving the three mechanisms which Barsalou claims are central to the formation of abstract concepts, namely, framing, selectivity, and introspective symbols. I have many experiences of my mother and form a simulator allowing me to recognize her and anyone exactly like her. But I am never

presented with evidence that there is anyone exactly like her, despite the thousands of opportunities for such a person to appear. Whenever I perceive my mother, there is never anyone else present with exactly her properties. The same can be said of the sun or the moon, which I also represent as cognized individuals.

This account accords well with Barsalou’s account of abstraction: “First, an abstract concept is framed against the background of a simulated event sequence” (sect. 3.4.2). In the case of my mother, the event sequence is drawn from all my experiences of her. “Second, selective attention highlights the core content of an abstract concept against its event background” (sect. 3.4.2). The selected core content of the abstract notion of a cognized individual is that this particular simulator is *sui generis*, that one has never encountered two perceived individuals together fitting this simulator. “Third, perceptual symbols for introspective states are central to the representation of abstract concepts” (sect. 3.4.2). An introspective state involved in the abstract notion of an individual is comparison of percepts.

Consider, briefly, classic cases of the identity relation as expressed by *Clark Kent is Superman* or *The Morning Star is the Evening Star*. Each such sentence seems to express a proposition equating two different cognized individuals. Before Lois Lane realized that Clark Kent is Superman, she had two distinct cognized individual simulators. She always and only saw Clark Kent wearing glasses and a baggy suit in the newspaper office; and she always and only saw Superman flying through the air in his red and blue cape and catsuit. On the day when she saw Clark Kent become Superman, these two individual concepts merged into a single, more complex cognized individual. If, at some later date, she actually explained to someone, in English, “You see, Clark Kent IS Superman,” she was not thereby reflecting two separate individuals in her own cognition, but collaboratively assuming that her hearer would not yet have merged the two concepts.

There is no space to explore the role that language plays in establishing abstract cognized individuals. The grammaticalization of deictic terms into definite determiners, such as English *the*, and the fossilization of definite descriptions into proper names, like *Baker* and *The Rockies*, have given us devices for expressing the abstract notion “individual.” The existence of stable cognized individuals in our minds, as opposed to merely transient perceived individuals, in terms of which Barsalou mainly writes, must be central to the ability of humans to “construct simulations jointly in response to language about nonpresent situations, thereby overcoming the present moment” (sect. 4.2, para. 3).

Identity, individuals, and the reidentification of particulars are classic problems in metaphysics and the philosophy of language. Barsalou has articulated a view of cognitive symbols which, if developed, can shed valuable light on these classic problems.

Creativity of metaphor in perceptual symbol systems

Bipin Indurkha

Department of Computer, Information, and Communication Sciences, Tokyo University of Agriculture and Technology, Tokyo 184-8588, Japan.

bipin@cc.tuat.ac.jp

Abstract: A metaphor can often create novel features in an object or a situation. This phenomenon has been particularly hard to account for using amodal symbol systems: although highlighting and downplaying can explain the shift of focus, it cannot explain how entirely new features can come about. We suggest here that the dynamism of perceptual symbol systems, particularly the notion of simulator, provides an elegant account of the creativity of metaphor. The elegance lies in the idea that the creation of new features by a metaphor proceeds in the same way as the creation of regular features in a perceptual symbol.

Barsalou has outlined an ambitious program to reaffirm the role of perception in cognition. The proposal is indeed quite sketchy, as

the author himself explicitly acknowledges, and many details will have to be worked out before a skeptic can be convinced of its viability. However, I find it intriguing, and applaud the author's attempt to rescue perception from the back seat and put it back in the driver's seat where it belongs. Indeed, the perceptual grounding of all cognition and the close bi-directional coupling between perception and cognition have largely been ignored by cognitive scientists and philosophers alike for most of this century. Chalmers et al. (1929) and Goodman (1976; 1978) are among the few notable exceptions to this general trend. For instance, Goodman echoed an essentially Kantian theme when he wrote: "Although conception without perception is merely *empty*, perception without conception is *blind* (totally inoperative)" (Goodman 1978, p. 6, emphasis Goodman's). Though this slogan is certainly witty, a comprehensive framework uniting conception and perception has not yet been attempted, as far as I am aware, until this target article by Barsalou.

I would like to focus my comments on a specific phenomenon having to do with the creativity of metaphor, where a metaphor creates a novel feature of an object or a situation. Recognizing the role of perception has a particularly illuminating effect in explaining this phenomenon, and we show how it can be addressed in the framework of perceptual symbol systems outlined by Barsalou.

Philosophers have long noted that metaphors can create new perspectives or new features on a certain object, situation, or experience; and in this creation lies the true cognitive force of metaphor (Black 1962a; 1979; Hausman 1984; 1989). This kind of creativity in problem solving has been studied, among others, by Gordon (1961; 1965), Koestler (1964) and Schön (1963; 1979). Recent psychological research has also demonstrated this aspect of creativity in understanding metaphorical juxtaposition in poetry (Gineste et al. 1997; Nueckles & Janetzko 1997; Tourangeau & Rips 1991).

Cognitive psychology and cognitive science research have largely addressed metaphor within amodal symbol systems, where metaphors are seen to arise from mappings between the symbols of two domains. The best these approaches can do is to use the notion of salience, together with highlighting and downplaying, to argue that creativity of metaphor consists in highlighting low-salience features of an object or situation. However, highlighting and downplaying does not explain how completely new features can be created.

An example can perhaps illustrate this point. In Stephen Spender's well-known poem "Seascape," the poet compares the ocean to a harp. In reading the poem, one's attention is invariably drawn to the way the rays of the sun reflect on the ripples of a calm ocean. This is not so much a matter of highlighting certain aspects of the amodal representation of the ocean, but creating a new representation for it. If one insists on explaining this as highlighting, then the amodal representation of the ocean quickly grows to an enormous proportion, as it must include every possible feature that can ever be associated with the ocean by any metaphor or any other cognitive mechanism. (See also Indurkha 1992.)

The notion of a simulator, however, allows for a rather elegant explanation of how new features can be created. As the simulators of both the harp and the ocean are activated, they try to build a simulation together. In this process, certain features of each cooperate (we might say that they resonate), while others cancel each other out. In this example, one might build an association between the images of the light reflecting on the strings of a harp as they lie waiting to be strummed and the sunlight playing on the ripples that seem to stand still in a calm ocean. The fact that the simulators contain the perceptual information about their referents is very crucial here, as the resonance occurs between the perceptual components, which cannot be reduced to a mapping between the symbolic levels. (After the metaphor has been assimilated, one can try to explain the metaphor verbally and symbolically. However, in many such metaphors, the verbal explanations are usually long and tortuous, and often hopelessly inadequate.)

The resonance of features might be unique in the sense that this particular confluence of features might never have existed before,

and could not have been activated by either of the concepts alone. Thus, a new feature emerges, which can then be explicitly incorporated in some perceptual symbol schema (see also Indurkha 1998). What is elegant about this account is that the emergence of new features by a metaphor is akin to the way features are first created, when a perceptual symbol schema is formed directly by selective attention (cf. sect. 2.2).

This account is akin to Fauconnier and Turner's conceptual blending approach (Turner & Fauconnier 1995; see also Fauconnier's accompanying commentary in this issue), where they show how concepts from many spaces blend together to produce metaphorical meanings. Barsalou's simulators extend Fauconnier and Turner's theory in a significant way by not only incorporating the perceptual dimension, but also giving it a paramount role to play; so that the account of creativity of metaphor outlined above can better be dubbed "perceptual blending."

Indeed, imagery and perceptual episodic memory have been known to play a key role in understanding certain metaphors (Marschark et al. 1983; Paivio 1979). This claim has been strengthened by recent neurolinguistic research (Bottini 1994; Burgess & Chiarello 1996), and some researchers argue that metaphors are essentially grounded in perception (Dent-Read & Szokolszky 1993). What Barsalou has tried to show us here is that metaphors are not special in this way, because all concepts are essentially grounded in perception.

The uncanny power of words

Paul J. M. Jorion

Théorie et Praxis, Maison des Sciences de l'Homme, 75270 Cedex 06, Paris.
 paul.jorion@email.msn.com aris.ss.uci.edu/~jorion

Abstract: In their quality as acoustic or visual percepts, words are linked to the emotional values of the state-of-affairs they evoke. This allows them to engender meanings capable of operating nearly entirely detached from percepts. Such a laying flat of meanings permits deliberation to take place within the window of consciousness. In such a theatre of the imagination, linguistically triggered, resides the originality of the human psyche.

I'm going for a walk in a wilderness park in Southern California. At the entrance of the park I am handed a leaflet about possible sightings of mountain lions. There are some instructions about how to act when faced with a puma. It sounds pretty exciting. I am on my own. I start walking, wondering how likely it is that I will be faced with the animal. Soon enough my attention drifts away, I'm day-dreaming while I walk, more or less absorbed in the "inner monologue" of the stream of consciousness. At some point I hear "within my head" – as part of inner speech: "I'm seeing one . . . I'm seeing one." This makes me focus on the scene. Indeed (the "indeed" with its element of pause is essential in the process), I suddenly realize that about fifty yards away from me on the trail, there's a puma, heading away from me into the brush.

What does this little piece of introspection reveal? There has been a visual perception of the animal no doubt, but it accedes to consciousness in a very indirect manner: the actual awareness is triggered by a sentence of inner speech, that is, the visual percept generates a piece of linguistic behaviour that then focuses my visual attention. I must have seen the animal *unconsciously*, which I then only properly see at the conscious level once the percept has been translated into words within inner speech.

A number of supposed coincidences work undoubtedly in the same manner. I am at the station, and I think to myself, "This good Jonathan, what a pity I haven't seen him for so many years." And lo and behold! Would you believe this? Here's Jonathan! What an extraordinary coincidence! Well no: I must have perceived him unconsciously (the mystery of peripheral vision?) half a second or so before. Enough time to have the inner speech elaborate on how true a friend he is.

Reinventing a broken wheel

Barbara Landau

Department of Psychology, University of Delaware, Newark, DE 19716.

blandau@udel.edu

This is what makes us human beings: we don't perceive hunger, we hear instead the inner voice saying "I kind of feel hungry." With us, the Word is indeed truly at the beginning. Although our perception is just like that of any other superior mammal: we, the speaking mammal, are faced with the concept before the percept (acute pain might be an exception).

What the word and the sentences in which they combine allow is deliberation: "Shall I pursue the walk toward the puma, or shall I retreat – just in case?" With any other mammal, there is probably very little deliberation, very little scrutinizing of alternative lines of conduct: The percepts elicit a combination of emotions that trigger the composition of a hormonal cocktail leading to a motor response, without any necessity for a display in the "window" of consciousness (Jorion 1999). What is amusing about Buridan's thought experiment about an ass paralyzed between two equidistant sacks of oats is precisely that donkeys do not act that way: you need reasoning with sentences to hesitate in this manner. Buridan's ass would go for the first sack that entered his visual field. The level where consciousness operates is essentially and crucially the level of the concept. And this is why – contrary to current popular views – there is probably no consciousness without the concept, the word, that is, no consciousness in animals.

Even so, with the example of the puma I have come a long way toward giving Barsalou's thesis some plausibility. In many other instances we are dealing with semantic issues pretty distant from pure perception. Identifying "perceptual symbols" with concepts as Barsalou does, breaks down somewhere along the road: of course there is some connection between words and percepts. Of course, saying "The apple, here on the table," has some relationship with perception: it is an instance of putting a percept into words. But take another example, which Barsalou will be at pains to dismiss as it is a proposition borrowed from his target article: "an amodal symbol system transduces a subset of a perceptual state into a completely new representation language that is inherently nonperceptual" (sect. 1.2). What proportion of the meaning of Barsalou's sentence actually boils down to percepts or even to "perceptual symbols"? Very little indeed: is ten percent a generous figure? Unless . . .

Unless one is prepared to regard concepts as what they likewise are: percepts, auditory in speech, visual in writing (and hallucinatory in the inner monologue?). But then the issue addressed by Barsalou dissolves into a quagmire of the universality of the percept. But here lies the solution to the question he aptly raises: a common feature underlies both percepts and concepts; their intrinsic association with an affective value, part of the affect dynamics which motivates and moves us along.

It is such a network of affective values linked to the satisfaction of our basic needs, and deposited as layered memories of appropriate and inappropriate responses, that allows "imagination" to unfold: to stage simulations of attempts at solution. What makes "The lion, the witch, and the wardrobe" an appealing title? That in the context of a child's world the threesome brings up similar emotional responses and are therefore conceptually linked (Jorion 1990, pp. 75–76).

It is part of the originality of human experience through the medium of language that the affective dynamics have the capacity to operate in a world of words only, dissociating themselves from any immediate physical resonance, that is, separated in an essential manner from the experience of sensual percepts (other than words). Barsalou's sentence quoted above is a convincing instance of such a process. "Can you open the window?" say I, and the window gets opened. The miracle of words. But also, "Tell me if proposition 209 is unconstitutional," and you may say Yes or No. Such is the power of language.

Abstract: Barsalou is right in arguing that perception has been unduly neglected in theories of concept formation. However, the theory he proposes is a weaker version of the classical empirical hypothesis about the relationship between sensation, perception, and concepts. It is weaker because it provides no principled basis for choosing the elementary components of perception. Furthermore, the proposed mechanism of concept formation, growth and development – simulation – is essentially equivalent to the notion of a concept, frame, or theory, and therefore inherits all the well-known problems inherent in these constructs. The theory of simulation does not provide a clearly better alternative to existing notions.

Barsalou is right in arguing that perception has indeed been unduly relegated to the back burner in theories of concept formation specifically, and cognition more generally. He is also right that the rich nature of perception has been ignored, and that it has often been mischaracterized as analogous to a record – a blurred version of what impinges on the sensory systems. Finally, he is right that a key problem in understanding the nature of concepts (especially lexical concepts) concerns how many of these are grounded in our perceptual experience.

However, the theory Barsalou proposes does not provide a strong and clear alternative to existing theories. Moreover, although he intends it to be a reinvention of the classical empiricist hypothesis, it is actually much weaker than the theories of Locke and Hume, and founders in its lack of specification of what simulation could be. It is weaker than the classical empiricist hypothesis because it provides no principled basis for choosing the elementary components of perception which will form the building blocks for concepts. The classical theory assumed that the atoms were those sensory impressions defined by the structure of the sense organs as they interact with the physics of the world. Operating on these were the processes of association and reflection. The pitfalls of constructing everyday concepts – such as Table, Ball, Above – from such elementary sensations are well known; but the advantage of such a system was its principled nature and its clear circumscription of what could be the elementary components of a compositional system (see Fodor 1981b).

Barsalou is right that much of recent history in cognition has been occupied with understanding the structure of those higher-level concepts that must play a role in interpreting perceptual features. Many have noted that the importance of any property or feature must be considered in the context of a particular concept in order to gain an appropriate weighting; this is why we all understand that feathers are somehow very important to the concept of a bird, but much less important (if important at all) to the concept of a hat. Such observations correctly suggest that such variable weighting demands a theory of the underlying concept – just what we set out to explain by listing features. The burden of explaining lexical concepts thus lies in understanding the concept itself – under what explanations the properties are related – and this is not likely to be forthcoming from a list of properties or features – even if these are expanded to include arguably nonprimitive features such as feathers.

Barsalou is also right that the emphasis on such higher-level concepts has often been made at the expense of understanding the relationship between what we now understand about perception, and these higher-level concepts. Nowhere is this imbalance more clearly shown than in the literature on lexical concept acquisition. A preponderance of recent research has been designed to show that young children can go beyond the limits of perception, making inferences about category membership even in cases where perceptual information is degraded or absent. These demonstrations clearly show top-down effects in young children; but they do not support the notion that perceptual organization is less than crucial in the early formation of these representations. Indeed, re-

cent work in the infancy literature provides compelling demonstrations of the extent to which perceptual organization during the first year of life can support the formation and use of categories (Quinn & Eimas 1996). Perceptual systems provide richly structured and dynamic information from the earliest developmental moments; and understanding the relationship between these systems and later modes of categorization is clearly a critical component of understanding human cognition.

However, the mechanism that Barsalou proposes is unlikely to move us forward in this understanding. Stimulation is essentially equivalent to the notion of a concept, frame, or theory. As such, it has all of the same virtues (being able to explain the dynamic, structured, and interpretive nature of concepts), and inherits all of the same flaws. Consider Barsalou's claim to have made progress in understanding how we represent Truth (sect. 3.4.3): a simulated event sequence frames the concept. But what is the concept? Simulations devoid of content do no better at characterizing our knowledge than concepts devoid of content. Without understanding the underlying notions of even such modest concepts as Bird, Hat, or Above, we are in no better position to understand how it is that we manage to gain, generalize, or ground our knowledge.

Latent Semantic Analysis (LSA), a disembodied learning machine, acquires human word meaning vicariously from language alone

Thomas K. Landauer

*Institute of Cognitive Science, University of Colorado at Boulder, Boulder, CO 80309. landauer@psych.colorado.edu
psych-www.colorado.edu/faculty/landauer.html*

Abstract: The hypothesis that perceptual mechanisms could have more representational and logical power than usually assumed is interesting and provocative, especially with regard to brain evolution. However, the importance of embodiment and grounding is exaggerated, and the implication that there is no highly abstract representation at all, and that human-like knowledge cannot be learned or represented without human bodies, is very doubtful. A machine-learning model, Latent Semantic Analysis (LSA) that closely mimics human word and passage meaning relations is offered as a counterexample.

Barsalou's hypothesis implies a path for evolution of human cognition that is consistent with its time course. Animals lacking specifically human brain structures make inferences, so the ability should exist in processes we share, primarily sensorimotor systems. But humans surpassed other animals by growing a huge volume of less specific cortex and developing far superior communication and abstract thinking. Language composed of largely perception-independent symbols is the chief vehicle. These more abstract processes communicate in both directions with perception, but part of our representational capacity must have attained a higher level of abstraction, autonomy, and conceptual scope.

Thus, I think Barsalou over-stretches what might be concluded. He apparently wants to persuade us that the components of cognition are all and only perceptual events or their simulations. Hence cognition is "embodied," always incorporates perceived idiosyncratic bodily states, is exclusively "grounded" in perception, represents nothing as abstract as discrete "amodal" symbols, and thus, he concludes, nonbiological machines are in principle incapable of human-like concepts.

If Barsalou were right, no machine lacking human sensory input could autonomously acquire human-like knowledge from experience. In particular, it could not learn to represent words and sentences, prime candidates for symbolic status, in such a way that their relations to each other would closely resemble any human's. However, there is an existence proof to the contrary. A machine-

learning implementation of Latent Semantic Analysis (LSA) (Landauer & Dumais 1997; Landauer et al. 1998a; 1998b) induces human-like meaning relations from a "senseless" stream of undefined arbitrary symbols. In these simulations, LSA is given the same naturally occurring text as that read by a typical human. The input is solely strings of letters separated into words and paragraphs. Without human assistance, LSA maps the initially meaningless words into a continuous high-dimensional "semantic space."

LSA passes multiple choice tests at college level, scores in the adult human range on vocabulary tests, mimics categorical word sorting and semantic relationship judgments, accurately mirrors word-to-word and sentence-to-word priming, correctly reflects word synonymy, antonymy, and polysemy, induces the similarity of phrases such as "concerned with social justice" and "feminist," and has been used – with no other source of data – to rank knowledge content of essays indistinguishably from expert humans. If word meaning for humans depended strongly on information derivable only from direct, embodied perception, this would be impossible.

The essential theory behind LSA is that the brain can use the mutual constraints inherent in the way events of any kind – including perceived objects and motoric plans – occur in each other's context over many millions of occasions to induce a complete mapping of the similarity of each to each. Of course, LSA, as currently implemented for word-learning, has limitations. Trivially, it has not lexicalized idioms and aphorisms and misses meanings that changed after the training corpus was collected. A more serious lack is a model for dynamically generated temporary meanings such as metaphorical expressions; LSA represents only statistically modal long-term meanings. (Note that Glenberg's unpublished studies cited by Barsalou as "strong evidence" against LSA are based on these sorts of incompleteness.)

LSA representations are not devoid of information about perception, emotion and intent. They derive from the way that living, feeling humans use words to describe experience. The way writers use "love" and "mauve" reflects much, if not everything, that experience has taught them. The point is that all this can be learned and carried by a program running in silicon. Were LSA also to have other sensory input than words, it would "ground" meanings in them too, but, given how well it does already, the difference would necessarily be small. The nonrealistic brittleness and unrealism of traditional AI "physical symbol systems" lies in their representational discreteness and their lack of meaning without respect human assignment of features, Barsalou's "transduction problem." LSA has no such problem. Language provides discrete symbols in words; LSA learns continuous valued meaning representations for them from linguistic experience.

LSA does not do everything that humans do with concepts, but enough to refute Barsalou's (and Searle's, etc.) strong assertions. It does not *have* perceptions and feelings, but it can do significant cognition *as if* it did. Boyle's law is not a gas, but it simulates some properties of gases well. LSA is not human, but it can simulate important aspects of human cognition. And, unlike simulating a gas, the practical consequence of simulating cognition is not just *like* cognition, it is cognition!

Barsalou's claim that a machine cannot have human concepts because each person has experience based on a unique body is a red herring. We need simulate only one typical human, similar, but not identical, to others. A critical aspect of human concepts is their mutual comprehensibility; the sharable properties are what we want to understand first and foremost. Likewise, the fuss over symbol grounding is overblown. The most primitive sensitivity to touch is useless without inference. There is nothing more physiologically or logically real about perception than about abstract cognition. Both are neural processes designed to promote adaptive function, only some of which needs to depend on perceptual input or feedback. It makes as much sense to speak of perception as grounded in inference as the other way around. Like LSA, humans can learn much that is needed for their inference and deduction engines from the shared use of symbols independent of percep-

tion, and sometimes do interesting abstract thinking, such as math, for which a strong tie to perception is either irrelevant or a nuisance.

A view from cognitive linguistics

Ronald W. Langacker

Department of Linguistics, University of California, San Diego, La Jolla, CA 92093-0108. rlangacker@ucsd.edu

Abstract: Barsalou's contribution converges with basic ideas and empirical findings of cognitive linguistics. They posit the same general architecture. The perceptual grounding of conceptual structure is a central tenet of cognitive linguistics. Our capacity to construe the same situation in alternate ways is fundamental to cognitive semantics, and numerous parallels are discernible between conceptual construal and visual perception. Grammar is meaningful, consisting of schematized patterns for the pairing of semantic and phonological structures. The meanings of grammatical elements reside primarily in the construal they impose on conceptual content. This view of linguistic structure appears to be compatible with Barsalou's proposals.

Barsalou's stimulating contribution converges in myriad respects with basic ideas and empirical findings of cognitive linguistics. Perceptual symbol systems are a promising vehicle for the implementation of cognitive linguistic descriptions, which in turn reveal a rich supply of phenomena that both support the general conception and offer significant challenges for its future development.

Perceptual symbol systems and cognitive linguistic approaches posit the same basic architecture, involving abstraction from experience, the flexible activation of abstracted symbols in top-down processing, and their recursive combination producing an open-ended array of complex conceptualizations. A central claim of "cognitive grammar" (Langacker 1987; 1990; 1991) is that all linguistic units arise from actual events of language use through the general processes of schematization and categorization. Grounded in bodily experience, elaborate conceptual systems are constructed through imaginative devices such as metaphorical projection (Lakoff & Johnson 1980; Lakoff & Turner 1989; Turner 1987), the creation of mental spaces (Fauconnier 1985; 1997), conceptual blending (Fauconnier & Sweetser 1996; Fauconnier & Turner 1998), and productive simulations. The same devices, as well as the conceptual structures already assembled, are dynamically employed in the construction of linguistic meanings. Rather than being mechanically derived from the meanings of its parts, an expression's meaning is actively constructed, largely in top-down fashion, from the fragmentary clues provided by its constitutive elements. Its coherence derives from their reconciliation in a complex simulation that draws on all available resources (conceptual, linguistic, and contextual).

The perceptual grounding of conceptual structure is a basic tenet of cognitive linguistics. "Image schemas" abstracted from sensory and kinesthetic experience have been claimed by Johnson (1987) and by Lakoff (1987) to be the building-blocks of conceptual structure, the abstract commonality preserved or imposed in metaphorical mappings (Lakoff 1990), and the structures most critically employed in reasoning. I myself have speculated that all conceptualization is ultimately derivable from certain irreducible realms of experience – including time, space, as well as sensory, emotive, and kinesthetic domains – however many levels of organization may intervene. I have further suggested that sensory and motor images are essential components of linguistic meanings (Langacker 1987). Talmy (1996) has coined the word "ception" to indicate the extensive parallels he has noted between *conception* and *perception*. The same parallels have led me (Langacker 1995) to use the term "viewing" for both (without however taking any position on the extent to which cognition might be visually grounded).

Fundamental to a conceptualist semantics is the recognition of our multifaceted ability to conceive and portray the same situation in alternate ways, resulting in subtly different linguistic meanings. This capacity for "construal" has numerous components that are posited through linguistic analysis and strongly supported by their descriptive utility. With striking consistency, these components of construal bear evident similarities to basic aspects of visual perception (Langacker 1993). For instance, linguistic expressions can characterize a situation at any desired level of precision and detail ("granularity"), as seen in lexical hierarchies such as *thing* > *creature* > *reptile* > *snake* > *rattlesnake* > *sidevinder*. This progression from highly schematic to increasingly more specific expressions seems quite analogous to the visual experience of seeing something with progressively greater acuity while walking up to it from a distance.

In visual perception, we can distinguish the maximal field of view, the general region of viewing attention (e.g., the stage, in watching a play), and the specific focus of attention within that region (e.g., a particular actor). An analogous set of constructs is required for the semantic description of linguistic expressions. The term *knee*, for example, evokes the conception of the body for its "maximal scope," selects the leg as its "immediate scope" (general locus of attention), within which it "profiles" the major joint (the conceptual referent, or focus of attention). In expressions that profile relationships, one participant – called the "trajector" – is accorded a kind of prominence analogous to that of the "figure" within a visual scene. For instance, *above* and *below* profile referentially identical relationships along the vertical axis; their semantic contrast resides in whether trajector status is conferred on the higher or the lower participant. Thus *X is above Y* is used to specify the location of X, while *Y is below X* specifies the location of Y. I take this contrast as being akin to the perceptual phenomenon of figure/ground reversal.

This is merely a sample of the parallels one can draw between visual perception and the constructs needed for conceptual semantic description. These constructs are not required just for expressions pertaining to visual, spatial, or even physical situations – they are fully general, figuring in the characterization of expressions relating to any realm of thought and experience, both concrete and abstract. Moreover, they also prove crucial for grammar. An expression's grammatical class, for example, is determined by the nature of its profile in particular (not its overall conceptual content), and trajector status is the basis for subjecthood (Langacker 1998). Once the importance of construal is fully recognized, the conceptual import of grammatical structure is discernible. In fact, the pivotal claim of cognitive grammar is that all valid grammatical constructs are meaningful, their semantic contributions residing primarily in the construal they impose on conceptual content. Lexicon and grammar are seen as forming a continuum fully describable as assemblies of "symbolic structures" (i.e., pairings between phonological and semantic structures). I believe this conception of linguistic structure to be broadly compatible with Barsalou's notion of perceptual symbol systems.

Can handicapped subjects use perceptual symbol systems?

F. Lowenthal

Cognitive Sciences, University of Mons, B-7000 Mons, Belgium.

francis.lowenthal@umh.ac.be www.umh.ac.be/~scoglab/index.html

Abstract: It is very tempting to try to reconcile perception and cognition and perceptual symbol systems may be a good way to achieve this; but is there actually a perception-cognition continuum? We offer several arguments for and against the existence of such a continuum and in favor of the choice of perceptual symbol systems. One of these arguments is purely theoretical, some are based on PET-scan observations and others are based on research with handicapped subjects who have communication

problems associated with cerebral lesions. These arguments suggest that modal perceptual symbols do indeed exist and that perception and cognition might have a common neuronal basis; but perceptual and cognitive activities require the activation of different neuronal structures.

1. Do the main differences between perceptual and nonperceptual symbol systems favor the perceptual ones? There seem to be two major differences between modal perceptual symbol systems and amodal nonperceptual symbol systems:

a. The modal “perceptual” symbols are global symbols, whereas the amodal “nonperceptual” symbol systems are combinations of small preexisting elements, such as words;

b. In the case of modal symbols, the coding procedure used to transform a perception into a symbol and store it in long term memory could be direct, from one neuronal zone to another; but in the case of amodal symbol systems, this coding procedure clearly requires the mediation of a preexistent, fully developed, symbol system such as a language.

This mediation of a preexistent, fully developed symbol system requires a framework, such as that described by Chomsky (1975), with its inborn LAD (language acquisition device), but it is totally unacceptable in a constructivist framework, such as those described by Piaget (Droz & Rahmy 1974), Bruner (1966), and Vygotsky (1962). The second difference between modal and amodal symbol systems thus favors perceptual modal symbols.

2. Do perceptually based activities favor conceptualization? We have observed and treated subjects with major brain lesions who have lost all form of language (Lowenthal 1992; Lowenthal & Saerens 1982; 1986) using a nonverbal technique (Lowenthal 1978) involving unfamiliar sets of objects that must be manipulated and under certain technical constraints that make certain actions possible and others impossible.

Lego bricks provide a simple example: When a subject is asked to make a path with Lego bricks on the usual baseboard, he discovers that the bricks used must be parallel to one side of the board. In the Lego world, there are only flat and right angles, no diagonals. Papert (1980) described how he discovered the first elements of algebra without instruction by playing with the gear wheels his grandfather had given him. Diñesí Attribute blocks, pegboards, and other devices such as geoboards can be used in the same way for more sophisticated exercises (Lowenthal 1986; 1992). All share a common property: they have constraints that make certain actions possible and others impossible. This provides the subject with a “logical structure” adapted to the device used. The technical constraints define the “axioms” of this logic. It has been suggested (Lowenthal 1986) that manipulating such devices helps the subject, first, to discover the relevant elements and, second, to combine small relevant elements to form bigger relevant blocks. This process continues until the subject reaches a satisfactory solution. It is hypothesized that the technical constraints on the devices enable subjects to formulate hypotheses in a nonverbal way, and to test and adapt them.

When our subjects were confronted with new problem situations, they began to respond by trial and error but quickly made significant progress. The concept formation was clearly based on manipulating and seeing. Hence the concepts were probably built at the Bruner’s first two (1966) representation levels.

Mauro (1990) has shown that subjects with clearly localized brain lesions can reconstruct a communication function even if the brain lesions are large and bilateral symmetrical; but this is not the case in subjects with diffuse brain lesions. These differences can be interpreted as follows: the nonverbal processing mentioned above enables subjects with large but localized brain lesions to use nondamaged neurons to reconstruct, at least partially, the equivalent of the damaged neuronal structures. This cannot be the case for patients with diffuse lesions because too many neuronal structures have been qualitatively damaged.

These observations imply an important role for the sensorimotor cortex as the substrate for the elaboration of concepts, provided there are still large “untouched” zones available. It may also

imply that, in the reorganization process, neurons could be involved in conceptualization even though they do not belong to the sensorimotor cortex. In the cases observed, we have enough data to conclude that the reconstructed communicative function is based entirely on the reconstructed symbols, and that these were created by the patient while he was using a technique based on manipulation and vision, hence perception. This does not allow us to conclude that the reorganized cognitive activity functions as an extension of perceptual activities.

3. Are there data showing that perceptual and cognitive activities have a common neuronal basis and share common neural structures? PET-scan data show that a majority of cortical neurons are involved in vision, which is a perceptual task; but they also suggest that the same neuron could be involved in different structures corresponding to different tasks. Mellet et al. have shown that the areas active in mental imagery tasks (images imagined but not actually perceived) are the same as the ones that are active during actual visual perception, except for the primary visual area V1 (Tzourio & Mellet 1998). They also showed that when concrete words are evoked (*banana*) some visual areas are active, but not V1, while no visual area is active in the case of abstract words (*justice*). Burbaud et al. observed that many prefrontal neurons are activated when subjects are engaged in computing tasks, such as addition. The active prefrontal neurons are not the same when the subject is engaged in a task such as “*add 624 and 257 mentally*” or in a task such as “*imagine you are writing on a blackboard and add 624 and 257*” (Burbaud 1998; Burbaud et al. 1995). According to these authors, neural activity differs depending on the subject’s preferred strategy and on his lateralization.

All these observations point to the same conclusion: the neurons of the sensorimotor cortex can play a role in cognitive activities but other neurons might also be involved. Cognitive mechanisms are not the same as perceptual ones: both mechanisms share some neural structures, but certain structures seem to be associated only with one. It might be the case that the activation of sensorimotor neurons is only associated with the evocation of symbols and concepts, but not with the cognitive treatment of these symbols.

Conclusions. There accordingly seem to be theoretical arguments in favor of perceptual symbol systems. Perceptually based activities also seem to favor the emergence of conceptualization in subjects with localized brain lesions, and this conceptualization seems to be the basis of a renewed cognitive activity; this does not seem to be the case in subjects with diffuse lesions. Nevertheless, PET-scan data seem to show that although perceptual and cognitive activities share common neural substrate, the structures involved can be very different. These findings argue against the existence of a perception-cognition continuum.

Whither structured representation?

Arthur B. Markman^a and Eric Dietrich^b

^aDepartment of Psychology, University of Texas, Austin, TX 78712; ^bProgram in Philosophy, Computers, and Cognitive Science, Binghamton University, Binghamton, NY 13902-6000.

markman@psy.utexas.edu dietrich@binghamton.edu
www.psy.utexas.edu/psy/faculty/markman/index.html
www.binghamton.edu/philosophy/home/faculty/index.html

Abstract: The perceptual symbol system view assumes that perceptual representations have a role-argument structure. A role-argument structure is often incorporated into amodal symbol systems in order to explain conceptual functions like abstraction and rule use. The power of perceptual symbol systems to support conceptual functions is likewise rooted in its use of structure. On Barsalou’s account, this capacity to use structure (in the form of frames) must be innate.

Barsalou’s perceptual symbol systems project is exciting and important. Among its virtues, it lays the groundwork for a more em-

bodied approach to cognition and it subtly changes the way cognitive scientists approach mental content, thereby changing what tools theories of mental content or intentionality can use.

A core distinction in cognitive science is the one between perceptual and conceptual representation. Barsalou points out that this distinction is problematic, because there is no successful demonstration of how these types of representations are connected. He suggests that this difficulty should be taken as a sign that perceptual representations are all there is. There are no amodal representations.

In order to make this account plausible, it is critical to demonstrate how key aspects of conceptual representation like abstraction can be cashed out in a perceptual symbol system (PSS). These functions of conceptual representation are accomplished in a PSS by assuming that perceptual representations have role-argument structure in the form of frames.

It is not surprising that Barsalou suggests that frames form the core of PSS. He has long regarded frame representations as being well suited to be the medium underlying conceptual representation (Barsalou 1992). Further, many cognitive scientists adopted representational systems with some kind of role-argument structure in order to facilitate the abstraction process (see e.g., Fodor 1981a; Markman 1999; Shank 1982). And finally, structured representations have been suggested as the basis of models of central perceptual processes like object recognition (Biederman 1987; Marr 1982). Thus, there is some reason to believe that perceptual representations are structured, and structured as Barsalou claims they are.

Naturally, the assumption that frames are critical to perceptual symbol systems leads to the question of where the capacity to represent information with frames comes from. There are two possibilities. One is that the frames develop over the life of a cognitive system from simple unstructured representations (e.g., vectors or independent features). A second possibility is that the capacity to build frames is an inherent capability of a PSS. These possibilities are just a version of the standard “learned versus innate” debate in cognitive science.

To our knowledge, all attempts to construct complex structured representational schemes starting with only unstructured representations have foundered. For example, attempts to account for the development of complex representational capacities using associative connectionist models were not successful (Fodor & Pylyshyn 1988; Marcus 1998). It is possible to create complex representational structures in a connectionist model, but such structure has to be built in ahead of time using other techniques that give the connectionist system a classical structuring capability (e.g., Shastri & Ajjanagadde 1993). A very good example of this is Smolensky’s interesting work on connectionism. His tensor product approach *does* have classical constituent structure already built in (see the series of papers: Smolensky 1990; Fodor & McLaughlin 1990; Smolensky 1995; and McLaughlin 1997).

In light of this difficulty, Barsalou assumes that the capacity to form structured representations is an inherent component of a PSS. That is, he assumes stimuli that contact a cognitive agent are converted into frame representations early on, and that the capacity to do this is innate in the system or organism. Although structured representations have been incorporated into models of perceptual processes, they are not a necessary component of models of perception (e.g., Ullman 1996).

The assumption that representations are structured is extraordinarily important for the PSS account. Once these frames are constructed, many of the techniques of concept formation and abstraction used by proponents of structured amodal representations can be incorporated into PSS including the ability to make similarity comparisons and to reason by analogy (Gentner 1983; Gentner & Markman 1997). Thus, once Barsalou assumes that representations in a PSS are frames, the ability of this system to account for higher level cognitive abilities is virtually assured.

Where the perceptual symbol system approach differs from previous approaches to structured representations is in assuming

that the components of representational frames are tied to perception rather than being derived from a central multimodal representation language (or perhaps from language ability itself). However, it is here that Barsalou’s approach has its greatest promissory note. It is a bold step to posit that the structure representations that form the basis of conceptual abilities are closely tied to perception. It is now critical to demonstrate how a true perceptual system could give rise to representations of this type.

Development, consciousness, and the perception/mental representation distinction

Lorraine McCune

Department of Educational Psychology, Rutgers University, New Brunswick, NJ 08903. mccune@rci.rutgers.edu

Abstract: Perceptual symbol systems provide a welcome alternative to amodal encapsulated means of cognitive processing. However, the relations between perceived reality and internal mentation require a more differentiated approach, reflecting both developmental differences between infant and adult experience and qualitative differences between consciously perceived and mentally represented contents. Neurological evidence suggests a developmental trajectory from initial perceptual states in infancy to a more differentiated consciousness from two years of age on. Children’s processing of and verbal expressions regarding motion events provides an example of the changing capacity for mental experience.

Barsalou defines a “symbol” as a “record,” a neurologically embodied result of the neural state underlying a given perceptual experience. Perceptual symbols so defined are described as the medium for all cognitive processing, both conscious and unconscious. The author is persuasive in his argument that an embodied modal system, directly relating biology to experience is a parsimonious and effective basis for mental processes. However, theoretical considerations regarding the term “symbolic” and developmental issues require further elucidation.

In the realm of consciousness, perception has traditionally been considered distinct from mental, “imaginal” or “symbolic” representation (Sartre 1984; Werner & Kaplan 1963). Perception is the ongoing experience of present reality, while representation, accomplished through the mediation of symbols, allows consideration of past, future, and counterfactual phenomena. Both types of conscious experience could have their basis in the perceptual symbol systems as defined, but the distinction between them needs to be addressed within the theory.

It has become common to roughly equate infant mentation with that of adults as Barsalou does, “From birth, and perhaps before, the infant simulates its expected experience and assesses whether these simulations map successfully into what actually occurs” (sect. 3.4.3, para. 6). The author acknowledges the importance of developmental transitions, but only up to a point, noting that “as infants’ perceptual and bodily systems develop, their conceptual systems should change accordingly. Nevertheless the same basic form of conceptual representation remains constant across both evolution and development, and a radically new form is not necessary” (sect. 4.2, para. 5). I would like to suggest that the infant’s initial capacity from birth and before is limited to direct perceptual experiences of which neural records may be kept, based on the infant’s attentional focus, as Barsalou proposes. But during the first two years the child develops the initial capacity for mentally representing reality, a transition which can be documented behaviorally as well as on the basis of changing neurological structures. Without such mental representation capacity the child is clearly incapable of entertaining propositions as described in section 3.2. The capacity for sensory and motor understanding of such relationships as “above” and “below” is not equivalent to propositional analysis.

Observable behavioral changes are demonstrated in activities

such as object search and play which are apparently under the child's conscious control. Evidence from postnatal neurological development favors an initial organization emphasizing sensory and motor areas which show rapid early synaptogenesis (e.g., visual cortex at 3–4 months; Huttenlocher 1994) with later growth in prefrontal areas leading to maximum density at about 12 months and increasing functional capacity as brain growth continues rapidly throughout the second year (Bates et al. 1994; Huttenlocher 1994; Vaughn & Kurtzberg 1992). This is the period of initial development and consolidation of mental representation as expressed in object search, representational play, and language. Changes in prefrontal functioning at 9–10 months differentiates 12-month-olds who tolerate long versus short delays on the Piagetian A not B search task, originally considered a measure of early mental representation (Fox & Bell 1993; Piaget 1954). The proposed inhibitory and integrational role of prefrontal cortex for this task and for the indirect object search task (Diamond et al. 1994) could also account for shifts in representational play associated with language changes in the second year of life (McCune 1995). High versus low language producers in the 14–20-month age range have exhibited different patterns of event-related potential (ERP) responses to comprehended versus familiar words, suggesting that the speech production itself may contribute to neurological organization for language processing (Mills 1994).

Sartre (1948) distinguished perceptual from imaginal (mentally represented) experience. A perceptual consciousness of a portrait attends to the pictorial properties of color, the appearance of the subject, and other information directly available to the senses. In a representational consciousness the portrait is merely a vehicle for evoking a mental experience of the individual portrayed. As adults we shift rapidly between these modalities in everyday life. There is reason to suppose the child achieves this capacity only gradually, the capacity to shift from perception to representation keeping pace with the developing capacity for representational experience. Both Piaget (1962) and Werner and Kaplan (1963) proposed bodily imitation of external events as the initial vehicle for internalizing perceptual and motor experiences which might later be represented in consciousness without such bodily accompaniment.

Motion events which infants witness and in which they participate provide a domain of experience including perceptual, motor, representational, and linguistic developments during the first two years of life which demonstrates the link between the earliest perceptual experience and the onset of propositional judgment. Objects and people in the world maintain their identity and appearance while undergoing a variety of motions in space. By 16 weeks-of-age, infants use kinematic information to discriminate the form of a pictured object from the background and to maintain the identity of a three-dimensional object, distinctions that cannot be made for stationary objects even at 32 weeks (Kellman 1993). Talmy (1985) proposed that “the motion situation” provides a structure in the real world that is consistently expressed syntactically and semantically across languages, but by varying linguistic means. Children beginning to speak produce words referring to common relational aspects of motion events which differ by part of speech across languages, while maintaining equivalent conceptual content. These words, first described by McCune-Nicolich (1981) as unified by a common underlying structure, give evidence of a universal set of cognitive/linguistic primitives consistent with their derivation from the child's experience of motion events.

Motion events display temporal event sequences, where the child's production of a relational word indicates a mental comparison of the present state of affairs to a prior, expected, or desired reverse alternative, for example: iteration or conjunction (*more, again*) and potential reversibility or negation (*no, ohoh, back*). These early single word “propositions” provide the first linguistic evidence that the child is using a system of mental representation – most likely a perceptual symbol system – for cognitive processing.

Simulations, simulators, amodality, and abstract terms

Robert W. Mitchell and Catherine A. Clement

Department of Psychology, Eastern Kentucky University, Richmond, KY 40475. psymitch@acs.eku.edu psyclement@acs.eku.edu

Abstract: Barsalou's interesting model might benefit from defining simulation and clarifying the implications of prior critiques for simulations (and not just for perceptual symbols). Contrary to claims, simulators (or frames) appear, in the limit, to be amodal. In addition, the account of abstract terms seems extremely limited.

Barsalou's proposal for grounding cognition in perceptual simulation is intriguing, and seems almost an elaboration of some ideas about the development of cognition from sensory-motor processes presented by Jean Piaget (e.g., 1945/1962; 1947/1972), who remains uncited. Barsalou attempts to avoid many of the problems of viewing cognition as based on perceptual “images” by letting in image-like structures (“simulations”) by the back door of neural processing. Although “perceptual symbols” are unconscious, componential, indeterminate, dynamic, and nonindividualistic records of neural states or activations related to perception, simulations can be conscious, determinate, static, and (somewhat) individualistic image-like entities, or the reverse (unconscious, indeterminate, etc.). Records of neural states/activation are themselves organized into a simulator (frame or concept, seemingly comparable to Piagetian schemata), which produces simulations based on perceptual states, thereby making clear that the simulation is not a copy of perception.

Barsalou's analysis is a useful way to avoid the presumed pitfalls of perception-based cognition but, oddly, “simulation” is never defined. Simulations are described as similar to “specific images of entities and events that go beyond particular entities and events experienced in the past” (sect. 2.4.2), with the important consequence that much of the enormous literature on perception-based cognition is still relevant to his model. It would seem, though, that many of the criticisms of the past (e.g., those of Berkeley and Hume) had not been directed, as Barsalou claims, to perceptual symbols, but rather to simulations. Thus, although the perceptual symbols avoid the problems about cognition derived from perception, it is unclear that simulations avoid these problems. In addition, in attempting to conceptually separate potentially conscious simulations from nonconscious perceptual symbols (sect. 2.1.1), Barsalou implies that consciousness is not at all necessary for cognition; but bringing up blindsight as an instance of non-conscious cognition ignores the fact that blindsight occurs only in people who previously had conscious visual experiences.

Simulations are repeatedly described as multimodal in the sense that they can incorporate multiple and diverse perceptual modalities, but the particular aspects described (frames, specializations) are always within one modality (predominantly vision) at a time. One potentially problematic aspect for the proposal is whether frames capable of being specialized multimodally are, in the limit, amodal when the same element in the frame can be represented in multiple modes. For example, problems arise when describing cross-modal knowledge such as the kinesthetic-visual matching present in mirror-self-recognition and bodily imitation (Mitchell 1994). If “body image” is represented both visually and kinesthetically, and both aspects can represent (or simulate) the same thing (the body), this representation/simulation must be either amodal (and accessible to both modes) or simultaneously multimodal. If it is simultaneously multimodal, the knowledge structure must be a connection (neural or other) between visually and kinesthetically stored simulations coordinated into a frame – a sensorimotor matching or mapping between frames, not just a possible specialization as either visual or kinesthetic. If this frame develops for the body image with all its possible specializations completely unspecified as to being kinesthetic or visual, isn't the frame then amodal? Similarly, rhythm can be detected in one's

own movements, in the visual experience (or memory) of another's movements, or auditorily in music. When one recognizes the identity in these different rhythms, is this a mapping among simulations, a mapping from simulations to perceptions, or an amodal detection of temporal consistency? More complex: How is a non-perceptual analogy recognized if a frame needs to be perceptually specialized?

Barsalou's analysis of truth, though impressive, seems more a description of how to decide if something is true while already having a notion of truth to decide about. A matching between a visual simulation and a visual experience could represent not only "truth," but also "similar," "comparable," and "looks like." Given that "truth" does not seem to be a perceptual construct, more seems necessary to justify the claim that people's core intuitive sense of "truth" is represented by matching simulation to perception. Similar problems seem likely with other perceptually unspecified terms such as "can," "might," "electricity," "ignorant," and even "thing." It will, then, be interesting to see how far perception-based cognition can go in elucidating how we understand the world and ourselves.

Introspection and the secret agent

Natika Newton¹

Nassau County Community College, New York, NY 11530.

nnewton@suffolk.lib.ny.us

Abstract: The notion of introspection is unparsimonious and unnecessary to explain the experiential grounding of our mentalistic concepts. Instead, we can look at subtle proprioceptive experiences, such as the experience of agency in planning motor acts, which may be explained in part by the phenomenon of collateral discharge or efference copy. Proprioceptive sensations experienced during perceptual and motor activity may account for everything that has traditionally been attributed to a special mental activity called "introspection."

Barsalou's target article presents a convincing case for perceptual rather than amodal symbol systems. Considered in light of a proposal like this one, the traditional amodal view appears, in hindsight, strikingly misguided and implausible. My only serious objection concerns Barsalou's remarks on introspection (sects. 2.3 and 2.4) which he sees as "an aspect of perceptual experience" like proprioception. He points out that introspective processing is relatively poorly understood.

It is questionable whether there is an aspect of perceptual experience that corresponds to "introspection." Barsalou does not define the term, and talks about it somewhat in the style of Locke who believed that we can attend to either outer or inner experience, and when doing the latter, will perceive mental operations – for example, comparing, representing, and so on. Why believe that there are specific brain mechanisms for processing perceptual experiences of other brain mechanisms? There would seem to be more parsimonious ways to explain what was traditionally called "introspection" without taking that risky step toward infinite regress, especially without evidence of any corresponding brain structures.

Barsalou does talk about selective attention and our ability to abstract, and suggests that in introspection, selective attention can focus on our "ability to represent something in its absence." This procedure seems to me to be a case of theorizing within a particular conceptual framework involving "the mind," "representation," and so on, but not requiring any ability to *experience* the "ability to represent." The concept of representing, if Barsalou is right, must be perceptually grounded like any other concept; but it can be grounded in the perceptual experience of external relations between representer and represented, not necessarily in the activity of representing all by itself. The ability to represent things may well be logically entailed by our representational experience

together with the theories we hold about the brain, but that entailment is not the same thing as the representing ability being an experiential aspect of the process of representation, in the way that, for example, visual features such as colors or outlines can be.

Instead of talking about introspection, I suggest that we look for subtle proprioceptive experiences that accompany activities like representing, attending to, and so forth, and that may be unnoticed but taken for granted as essential aspects of the experiences of those activities together with their sensory objects. For example, when one sees a visible object, one is aware not only of the visible features of the object, but of the proprioceptive sensations of focusing the eyes and orienting the body toward the object. These ordinary perceptual experiences may be all that is necessary to ground the concept of "seeing."

One aspect of experience that has traditionally seemed to require a special form of access like introspection is the *experience of agency* that accompanies all voluntary intentional activity, even subtle perceptual activity such as ocular control. But even this experience is proprioceptive. It appears to be produced by corollary discharge, or efference copy: information about motor commands to the muscles that is sent to the parietal lobe and to the cerebellum from the motor and premotor cortex, for purposes of comparison with the expected movement. This information about motor commands is attended to by the subject normally only when the expected feedback does not match the commands, or is lacking altogether as when the relevant limb is paralyzed. In ocular motor control, the experience is almost continuous, since the eyes must be constantly adjusted in response to shifts in the visual field (Houk et al. 1996). The subjective feeling of agency is strikingly absent when motor activity is initiated artificially in experimental conditions; in such cases the movement is "automatically referred to a source (an intention) distinct from the self" (Jeannerod 1997, p. 189). There is evidence that information from efference copy is held in working memory, and hence would be accessible to consciousness like any other perceptual experience (Houk et al. 1996).

Barsalou has supplied us with all the basic mechanisms needed to ground our entire conceptual system in sensorimotor experience. All that is needed is to tease out a few more of the tricks used by our bodies to inform us of what we are doing and what is happening to us so that we can be, or at least feel that we are, full intentional agents. Efference copy is one trick that bears more study. Unless systems like this prove to be inadequate to explain self-awareness, concepts like introspection should probably be kept in the archives.

NOTE

1. Author's mailing address is 15 Cedar Lane, Setauket, NY 11733.

Can metacognition be explained in terms of perceptual symbol systems?

Ruediger Oehlmann

University of Essex, The Data Archive, Psychology Unit, Wivenhoe Park, Colchester, CO4 3SQ, England. oehlmann@essex.ac.uk
dwww.essex.ac.uk/~roehl

Abstract: Barsalou's theory of perceptual symbol systems is considered from a metacognitive perspective. Two examples are discussed in terms of the proposed perceptual symbol theory. First, recent results in research on feeling-of-knowing judgement are used to argue for a representation of familiarity with input cues. This representation should support implicit memory. Second, the ability of maintaining a theory of other people's beliefs (theory of mind) is considered and it is suggested that a purely simulation-based view is insufficient to explain the available evidence. Both examples characterize areas where Barsalou's theory would benefit from additional detail.

Barsalou presents a series of very careful and challenging arguments for a perceptual theory of knowledge. This commentary will

consider Barsalou's view of introspection from the perspective of metacognition. A common view in various characterisations of metacognition is that the term describes mechanisms of monitoring and controlling cognition (Flavell 1976; Nelson & Narens 1990). This commentary will be centered around two examples that require thinking about an agent's knowledge and beliefs. The examples will be reviewed from the perspective of Barsalou's perceptual symbol system theory and difficulties to explain some metacognitive results in terms of this theory are highlighted. The first example involves feeling-of-knowing judgement.

Feeling of knowing (FoK) judgements have been considered as part of the monitoring process. The classical view is that these are judgements about whether a given currently nonrecallable item can be retrieved from memory (James 1890). A recent extension of this definition views FoK as a rapid, preretrieval stage (Miner & Reder 1994; Reder & Ritter 1992), and sees the function of FoK to be judging the expected retrievability of queried information. For example, Reder and Ritter (1992) presented their subjects with calculation tasks such as $23 \times 34 = ?$, and asked them to make a decision whether they wanted to obtain the solution by retrieval from memory or calculation. The results showed that subjects were able to make a very rapid initial evaluation of a question before an answer was attempted. Retrieving the actual answer required more time than the strategy choice decision, which was based on familiarity with problem parts rather than with the answer. Reder and Ritter therefore concluded that FoK judgement is based on familiarity with cues derived from the problem. Schunn et al.'s (1997) source activation confusion (SAC) model addressed the issue of cue familiarity by using input nodes to represent problem components such as 23, *, and 34. These nodes were linked to a node that represented the entire problem 23×34 that in turn was connected to a node which represented the solution. Activation passing from input nodes to the entire-problem node allowed FoK judgement whereas activation passing from this node to the solution node allowed answer retrieval.

Simulators could obviously be used to represent a problem and its solution. More specifically, the frames that form part of a simulator (target article, sect. 2.5, para. 1) could store the different components of the problem and solution as was indicated in Figure 3 of the target article (for a car). The solution could then be retrieved if the same problem was presented again. This approach would not support an intermediate representation of an input cue, however, as this is achieved in the SAC model. In contrast to FoK judgement that is essentially part of monitoring an agent's own knowledge, the second example is concerned with thinking about the beliefs of others.

Theory of mind (ToM) has been characterized as the ability to assign mental states to self and to others and to use knowledge about these states to predict and explain behavior (Leslie 1987; Premack & Woodruff 1978). A classical example of ToM is the false belief task described by Wimmer and Perner (1983). It involves the child Maxi, who puts chocolate in a blue cupboard and goes playing. In the meantime, the mother uses part of the chocolate and puts the remaining part into the green cupboard. Wimmer and Perner reported that when asked where Maxi would look for the chocolate after returning from playing, most children under four years of age attributed to Maxi what they believed themselves, that the chocolate is now in the green cupboard.

Barsalou relates introspection to representational and emotional states (target article, sect. 2.3.1, para. 1). A clarification is needed of how his theory relates mental states such as beliefs and desires to introspection. If these were meta-level experiences represented by simulators, it would mean that both beliefs and desires could be simulated. However, Gopnik and Slaughter (1991) demonstrated that three-year-olds can report mental states such as desires and perceptions that are just past but now different, whereas they confirmed that three-year-olds failed false belief tasks. In addition, three-year-olds pass discrepant belief tasks, which require that another person maintain a belief that is not misrepresented but different from the agent's belief (Wellman &

Bartsch 1988). Gopnik and Wellman explain the different results for false belief tasks and discrepant belief tasks by pointing out that the child fails to realize that the belief in the false belief task is misrepresented. If beliefs were simply simulated, these different results would not be understandable. Another explanation has been put forward that is based on the assumption that children's early conceptions of the mind are also implicit theories and that changes of these conceptions are theory changes (Gopnik & Wellman 1994; Leslie 1994; Leslie & German 1995). This assumption has been referred to as theory-theory; it argues that young children do not have an explanatory psychological construct that is based on a theoretical conception of belief. Therefore their explanations do not consider beliefs and especially false beliefs or misrepresentations. This suggests that rather than relying solely on simulation, a perceptual theory of knowledge should account for constructs derived from implicit theories as well.

In conclusion, two examples have been presented to indicate why it is difficult to explain some metacognitive results in terms of perceptual symbol systems. The first focused on frames and their organization in a simulator. In the SAC model (Schunn et al. 1997), there is clearly a functional distinction between units that represent input components and intermediate units that are involved in FoK judgement. Nodes that store solutions can only be retrieved *after* activation has been passed to relevant intermediate nodes. It appears that in Barsalou's perceptual symbol system, such a distinction would have to be clarified for links between frames before simulators would be able to support FoK judgement. The second example questioned simulations as the only approach to understanding beliefs and suggested that some results could be better explained by assuming implicit theories. Hence it might be useful if a theory of perceptual symbol systems allowed for explanations based on implicit theories as well as simulations.

Selecting is not abstracting

Stellan Ohlsson

Department of Psychology, University of Illinois at Chicago, Chicago, IL 60607. stellan@uic.edu

Abstract: Barsalou's hypothesis that mental representations are constructed by selecting parts of percepts encounters the same difficulties as other empiricist theories: They cannot explain concepts for which instances do not share perceptible features (e.g., *furniture*) or for which there are no relevant percepts (e.g., *the end of time*). Barsalou's attempt to reduce falsity to failed pattern matching is an elementary error, and the generativity of his simulators cannot be attained without nonterminal symbols. There is not now, and there never was, any reason to be interested in empiricist theories of knowledge. Abstraction is a fundamental aspect of human cognition.

Abstraction. Empiricist theories of knowledge claim that people derive knowledge from experience, either by extracting shared features from sets of exemplars (Aristotle, Locke) or by combining sense data (Carnap, Russell). Barsalou proposes a different process: to select and recombine parts of percepts.

Classical empiricist theories could not explain abstract concepts. Barsalou's theory is no more successful, because selection is the wrong tool for the job. A part of a thing is not more abstract than the thing itself. For example, a wheel is as concrete an entity as a car. Hence, the percept of a wheel is not at a higher level of abstraction than a percept of a car. Selecting is not abstracting.

A closely related difficulty for Barsalou's theory is that the instances of some concepts do not share any perceptible features. Consider *furniture*, *tools*, and *energy sources*. No perceptible feature recurs across all instances of either of these categories. Hence, those concepts cannot be represented by combining parts of past percepts.

Scientific concepts illustrate the same point (Ohlsson & Lehtinen 1997). The physicist's concept of a *mechanical system* includes

moons and pendulums. However, my brain cannot take parts of my percept of the moon and build a perceptual representation of a pendulum out of them. The path of the moon across the sky does not look at all like the swings of a pendulum, so the corresponding percepts do not share any part that can be used to encode the concept.

A slightly different challenge is posed by entities like *the health care system*, *the middle class*, *the recent election*, and *the Internet*. They lack spatio-temporal boundaries and have indeterminate physical embodiments, so they are not perceived in any normal sense. Hence, they cannot be represented by parts of percepts.

An even stronger challenge is posed by concepts that do not correspond to any percepts at all. Consider *the future*, *the end of time*, and *life after death*. I have never perceived these things, so how could my mental representations of them consist of parts of percepts? Mathematical concepts make the same point. What parts of which percepts are supposed to constitute my concepts of *the square root of minus one*, *infinite sets*, and *the derivative of a function*? Finally, consider fictional concepts. What parts of which percepts are supposed to constitute my concepts of *time travel*, *instantaneous matter transmitters*, or *telepathy*?

In section 3.4.2, Barsalou postulates two capabilities other than selection for constructing abstract concepts: the ability to represent event sequences and the ability to store percepts of introspective states. As the above examples show, abstract concepts are not typically associated with any particular event sequences or introspections. Barsalou's treatment of abstraction is entirely inadequate.

Truth and falsity. Barsalou proposes to reduce truth and falsity to pattern matching. This approach seems to work as long as we only consider positive matches. One might argue that to establish the truth of an assertion like *the cat is on the mat*, it is sufficient to look in the relevant direction and match the pattern ON(CAT, MAT) to the incoming perceptual information. If the pattern fits, then it is true that the cat is on the mat.

However, as implementers of the early production system languages quickly discovered, *falsity cannot be reduced to a failed pattern match*. Suppose somebody quietly pulls the rug out of the room while I am looking out the window and that the cat remains seated, enjoying the ride. Next time I look around, the pattern ON(CAT, MAT) no longer matches my perception, but I cannot conclude from this that *the cat is on the mat* is false. This is a computational formulation of the old insight that the absence of evidence is not evidence for an absence. For example, the fact that astronomers have so far failed to find any patterns in the sky that indicate communication attempts by extra-terrestrials does not prove that *there are aliens out there* is false. To propose to reduce falsity to failed pattern matching is an elementary technical error.

Generativity. Barsalou assigns wonderful powers to his simulators, but he does not describe how they are supposed to operate. However, it appears that simulators can generate infinitely varied representations ("the simulator for *chair* can simulate many different chairs under many different circumstances," sect. 2.4.3, para. 2).

This claim is incompatible with the perceptual symbol idea, because a process that can generate a range of representations cannot itself be one of those representations. For example, consider the sequence of measurements 15, 17, 19, . . . and the generative mechanism $N(i + 1) = f[N(i)] = N(i) + 2$. The function f is not itself a member of the sequence 15, 17, 19, . . . This observation is general: a generative mechanism requires nonterminal symbols (f , N , $+$, etc.). Hence, Barsalou's simulators, if they are to carry out the generative function he assigns them, cannot themselves consist of nothing but perceptual symbols.

Conclusion. There is not now, and there never was, any positive reason to be interested in empiricist theories of cognition. Empiricists are forever arguing, as does Barsalou, to survive the obvious objections. Better be done with such theories once and for all and focus on the fundamental and remarkable fact that humans are capable of forming abstractions. The most natural ex-

planation for this fact is that abstract concepts are exactly what they appear to be: internally generated patterns that are applied to perceptual experience to interpret the latter (Ohlsson 1993). A successful theory of human cognition should start from the principle that although knowledge is *of* experience it does not derive *from* experience.

A little mechanism can go a long way

David A. Schwartz,^a Mark Weaver,^b and Stephen Kaplan,^c

^aDepartment of Psychology, University of Michigan, Ann Arbor, MI 48109-1109; ^bCorvus Development, Inc., Ann Arbor, MI 48104; ^cDepartment of Electrical Engineering and Computer Science and Department of Psychology, University of Michigan, Ann Arbor, MI 48109-1109.
hebb@umich.edu weaver@corvusdev.com skap@umich.edu

Abstract: We propose a way in which Barsalou could strengthen his position and at the same time make a considerable dent in the category/abstraction problem (that he suggests remains unsolved). There exists a class of connectionist models that solves this problem parsimoniously and provides a mechanistic underpinning for the promising high-level architecture he proposes.

No cognitive theory should be required to answer all the questions we might ask about the origins and structure of knowledge (at least not all at once, in a single article), and Barsalou, to his credit, tells us explicitly which questions his theory does not attempt to answer. However, by restricting his attention to the high-level architecture (leaving out any discussion of detailed mechanisms) and by claiming that a laundry list of issues related to categories and concepts remain "unresolved in all theories of knowledge," Barsalou may be creating a misleading impression about the current state of perceptually-grounded theories of cognition – as well as failing to take advantage of some valuable sources of support.

Barsalou takes some pains to show us that perceptual theories of thought have a long and distinguished pedigree. We find it curious, therefore, that he neglected the history of perceptually-grounded connectionist theory. Barsalou cites Pulvermüller (1999) several times as an example of connectionist work that is compatible with a perceptual symbol-system perspective. But Pulvermüller's approach is directly inspired by Hebb's cell-assembly theory, which was, of course, proposed in 1949 in "The Organization of Behavior." We agree with Barsalou that not all connectionist models can be considered to support "perceptual symbols"; similarly, not all connectionist models can properly be considered as descendants of Hebb's cell-assembly proposal. However, there is a substantial group which can be so considered, a group that we have previously referred to as "active symbol" connectionists (Kaplan et al. 1990), which includes not only Pulvermüller (1999), but also, for example, Braitenberg (1978; [see also Braitenberg et al. "The Detection and Generation of Sequences as a Key to Cerebellar Function" *BBS* 20(2) 1997]), Edelman (1987), Grossberg (1987; [see also Grossberg: "Quantized Geometry of Visual Space" *BBS* 6(4) 1983]), Kaplan et al. (1991), and Palm (1982). Although Hebb's main concern was a cell-assembly's capacity to support persistence, not perceptual groundedness, cell-assemblies, as Hebb described them, do have the kind of perceptual groundedness that is critical for Barsalou. [See also Amit: "The Hebbian Paradigm Reintegrated" *BBS* 18(4) 1995.]

Barsalou cites the Hebbian strengthening of connections between active neurons as one way in which perceptual symbols might form, but he seems not to have considered that this same associative mechanism by itself ensures that the resulting conceptual structures will be abstract, schematic, and categorical. By the Hebbian learning process, cells that initially exhibit patterns of activity that are weakly positively correlated grow increasingly interconnected and come to function as a single unit; and cells that initially display uncorrelated or weakly negatively correlated firing

patterns will grow increasingly distinct functionally and end up as members of different neural circuits.

What kind of neural representation will this sort of process generate? Consider a layer of cortex whose neurons respond to different feature analyzers at the sensory interface (e.g., shape, color, texture). A first encounter with a given object would set up a fairly diffuse pattern of cortical activity and would lead to a modest strengthening of the connections among the active elements. Because every object is encountered in some context, some of the active elements would represent features of the object, whereas others would represent contextual features. Given repeated encounters with the same object in varying contexts, the neurons representing the object's features will grow into a highly interconnected network while those representing contextual features will tend to fall away. Hebb coined the terms "recruitment" and "fractionation" to describe these two processes whereby neurons become incorporated into or excluded from a network.

The same processes of recruitment and fractionation occurring at a stage further removed from the sensory interface can grow neural representations that stand for classes of objects. Given repeated encounters with a set of objects that share certain features, birds for example, the neural units responding to the most invariant features (e.g., feathers and beaks) will grow into a highly interconnected functional unit, whereas the more variable features (e.g., color, size) will be excluded from the set of core elements. There is in principle no limit to the number of different levels of abstraction at which this process might operate. Given sufficient cortical "depth" and sufficiently varied experience, we might expect an individual to develop neural representations that stand not only for classes of objects, but for classes of classes, and so on ad infinitum. An assembly of neurons that forms in this fashion will exhibit many of the properties Barsalou attributes to perceptual symbols. It will be schematic, in that it represents only a subset of the features that any actual object manifests at any given time. It subserves categorization, in that the same assembly responds to varying instances of some class of objects that have features in common. It is inherently perceptual, dynamic, and can participate in reflective thought. In addition, it allows for recognition on the basis of partial input, typicality effects, prototype formation, and imagery. And note that we get all of this, in a sense, for free. The mechanisms of recruitment and fractionation that build the required perceptual symbols at the same time isolate the invariant features of some class of objects without the need for any outside intelligence to guide the process along.

By contrast, Barsalou seems to place much of the burden for perceptual symbol formation on an unexplained process of selective attention. According to his theory, attention selects some partial aspect of a perceptual experience and stores this aspect in long-term memory. Moreover, he offers no account of why in any given perceptual experience attention might focus on some features and not on others. This we find troubling, for in the absence of any principled constraint on which subsets of experience come to serve as dynamic internal representations, Barsalou's perceptual symbols seem only marginally less arbitrary than the symbols postulated by the amodal and connectionist theories he rightly criticizes.

By distancing himself from connectionist approaches in general, Barsalou not only undermines his capacity to deal with the particular problem of categories and abstractions, but also his capacity to specify any sort of mechanism for the perceptual approach. He perhaps did not realize that some forms of connectionism, in particular the Hebbian or "active symbol" version, are well suited to support and give structure to his position.

Truth and intra-personal concept stability

Mark Siebel

Research Group in Communication and Understanding, University of Leipzig, 04109 Leipzig, Germany. siebel@rz.uni-leipzig.de

Abstract: I criticize three claims concerning simulators: (1) That a simulator provides the best-fitting simulation of the perceptual impression one has of an object does not guarantee, pace Barsalou, that the object belongs to the simulator's category. (2) The people described by Barsalou do not acquire a concept of truth because they are not sensitive about the potential inadequacy of their sense impressions. (3) Simulator update prevents Barsalou's way of individuating concepts (i.e., identifying them with simulators) from solving the problem of intra-personal concept stability because to update a simulator is to change its content, and concepts with different contents are distinct.

Barsalou's theory involves three claims about simulators that I wish to criticize. They concern (1) true propositions, (2) the development of a concept of truth, and (3) intra-personal concept stability.

1. According to Barsalou (sects. 2.4.4 and 3.2.3), one entertains a true proposition to the effect that a perceived individual belongs to a certain category if one's simulator for that category provides the best-fitting simulation of the perceptual impression one has of the individual. However, that account rests on a confusion between what seems to be true (to a person) and what is true.

Consider someone who sees a horse that looks like a donkey to him; the horse causes a neural representation that resembles the impressions he normally has when he sees a donkey. Consequently, the simulator that provides the best-fitting simulation of the person's impression is his simulator for donkeys, not for horses. Thus, according to Barsalou, the person does not only classify the horse as being a donkey, his conceptual system also provides a true proposition. Because the simulator for donkeys produces a satisfactory simulation of the horse, the horse should belong to the simulator's category. But, obviously, a horse does not become a donkey because it *seems* to be a donkey. Otherwise it would be nearly impossible to make a mistake in classifying an object. The resemblance between a perceptual representation of an object and representations created by a simulator for a certain category does not guarantee the object's category membership. Whether the object belongs to the category depends on the veridicality of those representations.

2. Barsalou's account of the way we acquire a concept of truth is infected by a similar problem. In Barsalou's view, people acquire a concept of truth by "learn[ing] to simulate the experience of successfully mapping an internal simulation into a perceived scene" (sect. 3.4.3, para. 2). They develop such a concept by, for example, comparing perceptual simulations they form by hearing the utterance of a sentence and perceptual impressions they have of a corresponding scene. If the simulation resembles the impression to a certain extent, they learn to reply to the utterance with "That's true."

Roughly, truth is correspondence to the facts: a representation is true if and only if it represents a realized state of affairs. What Barsalou's people learn, however, is not construing truth as correspondence to the *facts*, but as correspondence to certain *representations*, namely, sense impressions of perceived scenes. Their concept of "truth" is instantiated whenever the simulation caused by hearing an utterance corresponds to the impression caused by looking at a certain scene – regardless of whether the impression itself is veridical. Barsalou's people are not sensitive about the potential inadequacy of their sense impressions. In a word, they do not develop a concept of truth because it is too easy to elicit assent from them: they say "That's true" whenever the way in which they conceptualize the meaning of the utterance resembles the way in which they perceive the crucial scene. They do not know what it is for an utterance to be true because they do not take into consideration that their perceptual impressions can lead them astray.

3. As Barsalou (sect. 2.4.5) himself recognizes, an adequate

theory of concepts must allow for a certain extent of intrapersonal concept stability. It must allow, for example, that the concept one uses in classifying an object on Thursday is the same as the concept one used on Monday to classify another object. Barsalou (sects. 2.4.3–2.4.5) tries to achieve stability by identifying concepts with simulators: even if different conceptualizations (i.e., simulations) are responsible for classifying different objects as being members of the same category, the same concept is used in these processes as long as the same simulator provides these best-fitting simulations.

That way of individuating concepts, however, remains too fine-grained to save intra-personal concept stability, for Barsalou (sects. 2.4.4 and 3.2.2) says that every successful simulation updates the corresponding simulator by integrating perceptual symbols of the categorized object into it. In other words, the categorization of an object supplies the simulator with new perceptual information about the members of its category so that it is able to create further simulations afterwards. Thus, updating a simulator (i.e., a concept) means to change its content; and since concepts with different contents are distinct, the update leads to another concept.

Concepts are to be individuated by their *extension* (their referents) and their *content* (the information they contain about their referents). Bertrand's concept of donkeys is already different from Susan's concept of horses because it represents another category. And if Bertrand's concept of horses does not contain the same information as Susan's concept of horses, they are likewise distinct, although they have the same extension. To use a Fregean term, even concepts for the same category differ if their "modes of presentation" are distinct, that is, they represent that category in different ways.

Hence, an update of a simulator entails that it is no longer the same concept. Although the category of objects it represents remains the same, we have a different concept afterwards, because there is something added to its content, namely, perceptual information about the object categorized before. Hence it is hard to see how Barsalou's account could save concept stability within individuals. If the categorization of an object is followed by an update of the corresponding simulator, then the next assignment of an object to the same category does not arise from the same simulator because its content changed. Because different contents imply different concepts, updating prevents a person from using the same concept twice.

In order to save concept stability within individuals, one must not identify concepts with simulators as a whole, but only with a certain *core* of them. Much as shared concepts between individuals only require *similar* simulators, the identity of a concept within an individual merely requires that the underlying simulator not have changed *too much*.

A perceptual theory of knowledge: Specifying some details

Aaro Toomela

Department of Psychology, University of Tartu, Tartu, Estonia EE 50410.
toomela@psych.ut.ee

Abstract: We attempt to resolve some details of Barsalou's theory. (1) The mechanism that guides selection of perceptual information may be the efferent control of activity. (2) Information about a world that is not accessible to the senses can be constructed in the process of semiotic mediation. (3) Introspection may not be a kind of perception; rather, semiotically mediated information processing might be necessary for the emergence of introspection.

Barsalou describes a perceptual theory of knowledge for which the need emerges from the historical divergence of the fields of cognition and perception in psychology. Russian psychology, how-

ever, especially Vygotsky-Luria's school of psychology, never separated thinking from perception (Luria 1969; 1973; 1979; Vygotsky 1934/1996). Barsalou also suggests that there is no theory of knowledge to explain when emergent features arise and why (note 18). Yet the primary objective of the field of semiotics is to understand semiosis and the knowledge-making activity (Danesi 1998). The question of how novel information is generated is thoroughly analyzed in the Tartu-Moscow school of semiotics (e.g., Lotman 1983; 1984; 1989). Thus, there are some ideas in Vygotsky-Luria's psychology and in semiotics that may be useful for filling gaps in Barsalou's theory.

Selection of information may be guided by the efferent control of activity. According to Barsalou, selective attention has a central role in the construction of perceptual symbols (e.g., sect. 2.2), but he does not explain how the cognitive system knows where to focus attention during the symbol formation process (note 7). He assumes that a conceptual system may be based on perceptual symbols. However, if we do not know how attention selects perceptual information then there is a possibility that the cognitive system that underlies selective attention is amodal. (Of course that system may also be modal, but the point is that Barsalou cannot exclude amodal systems from his theory without demonstrating that such a selection mechanism – which "knows" where to focus attention – is modal.)

Luria's theory offers a modal solution to the problem of attention in the form of a special system in cognitive functioning that is responsible for programming, regulating, and monitoring activity. In addition to afferent modalities, there is an efferent (basically motor) modality. The same system supports attention (Luria 1969; 1973). It is possible that isolation of novel information in perception for subsequent storage in long-term memory is realized by that efferent system. There is some supporting evidence, for example, the disappearance of stabilized retinal images (Alpern 1972; Cornsweet 1970), showing that activity (eye-movement) is necessary for visual sensation. There is also evidence that visual agnosias are accompanied with disorders in eye-movements (Luria 1969; 1973). This finding indicates that activity is also necessary for the functioning of central perceptual mechanisms.

The relationship between consciousness and introspection is unclear. Barsalou suggests that conscious experience may be necessary for the symbol formation process to occur initially (sect. 2.2). By "consciousness," Barsalou means awareness of experience or qualia (personal communication, December 1, 1998). At the same time, he assumes that the term "perceptual" refers both to the sensory modalities and to introspection (sect. 2.3.1). According to him, introspection represents experiences (representational states, cognitive operations, and emotional states; sect. 2.3.1). It is not clear how it is possible to be aware of experiences for building perceptual-introspective knowledge about the same experiences. We return to this problem below.

Information about a world that is not accessible to the senses cannot be based on perceptual symbols alone. Barsalou assumes that every kind of knowledge can be represented with perceptual symbols. The perceptual theory of knowledge is not very convincing regarding abstract concepts even though the solution proposed there is not entirely implausible. However, there exists a kind of knowledge that must be amodal in essence: the knowledge about a world which is qualitatively out of reach of our senses. Humans do not possess perceptual mechanisms for perceiving electromagnetic fields, electrons, and other submolecular particles. None of the modalities of the human perception can react to a neutrino. How such knowledge is constructed is not explained in Barsalou's theory (note 18 partly recognizes this problem).

One solution to the question of how novel information is generated is proposed by Lotman (e.g., 1983; 1984; 1989). According to him novel information emerges only if at least two different mechanisms of information processing interact, the novelty emerging when the results of one type of information processing is "translated" into the other. Humans possess such mechanisms, sensory-perceptual and amodal-linguistic. Novel information is

created in the process of semiotic mediation, in the “dialogue” between the perceptual and linguistic mechanisms. In the process of semiotic mediation, perceptual events acquire entirely new meanings (cf. Lotman 1983; Toomela 1996). The same basic mechanism may be responsible for the emergence of novel knowledge about a world that is accessible to the senses. In this case novelty emerges when information about the same object or phenomenon is processed in a dialogue between different sensory modalities.

Introspection may be a kind of semiotically mediated knowledge. It is not possible to go into detail here, but it is worth considering the possibility that introspection and consciousness represent a qualitatively novel knowledge that is not accessible to our senses. It might not be accidental that “introspective processing is poorly understood” (sect. 2.3.1). To the best of my knowledge, there is no direct evidence that humans possess an introspective system analogous to sensory mechanisms. Rather, introspection might be a result of semiotic mediation (cf. Luria 1973; 1979; Vygotsky 1925/1982; 1934/1996).

The idea that humans possess different mechanisms for processing the same information fits with some of Barsalou’s conclusions. First, the right hemisphere represents individuals in perception and comprehension and the left hemisphere represents simulators that construe them (note 29). It is also possible, however, that right hemisphere processes information more in terms of perceptual characteristics and the left hemisphere more with “linguistic simulators” (cf. Lotman 1983). Second, in sect. 3.4.2, Barsalou refers to the finding that abstract concepts tend to activate frontal brain regions. Prefrontal brain regions are also known to be related to verbal regulation of behavior and verbal thinking (Luria 1969; 1973). The construction of abstract concepts may require semiotically mediated thinking supported by frontal structures. Finally, Barsalou suggests that the dissociation of visual agnosia and optic aphasia may provide evidence for two different ways of activating a common representational system (note 3). That interpretation also accords with the idea that the same information might be processed in different ways in the same cognitive system.

Most of my remarks have concerned information in footnotes where undeveloped aspects of the theory are recognized. There are many theoretically important ideas in the body of the target article that not only fit with other perceptual theories but also extend our understanding of the role of perception in cognition. Some interesting ways to improve the perceptual theory of knowledge might be found in Vygotsky-Luria’s school of psychology and in Lotman’s school of semiotics.

ACKNOWLEDGMENT

This work was supported by The Estonian Science Foundation.

External symbols are a better bet than perceptual symbols

A. J. Wells

Department of Social Psychology, The London School of Economics and Political Science, London WC2A 2AE, England. a.j.wells@lse.ac.uk

Abstract: Barsalou’s theory rightly emphasizes the perceptual basis of cognition. However, the perceptual symbols that he proposes seem ill suited to carry the representational burden entailed by the architecture in which they function, given that Barsalou accepts the requirement for productivity. A more radical proposal is needed in which symbols are largely external to the cognizer and linked to internal states via perception.

I strongly endorse Barsalou’s emphasis on the perceptual nature of cognition and I share his view that theories based on amodal symbol systems suffer from serious difficulties. Transduction and symbol grounding (cf. sect. 1.2.2), in particular, are fundamental problems. Barsalou proposes that at least some of the problems

can be solved by hypothesizing a class of internal representations, that is, perceptual symbols, that stand in a different relation to the proximal stimuli that produced them than do amodal symbols. The proposal looks radical at first sight but on further examination seems rather modest. It is precisely the modesty of the proposal that gets it into difficulties and I wish to recommend a more radical proposal.

Much traditional theorizing about the linkages between the external world and internal thought processes acknowledges the following broad classes of entities. There are distal events and objects, proximal stimuli, perceptual states and cognitive states. Barsalou seems to acknowledge these classes. Roughly, it is assumed that distal entities cause proximal stimuli, which cause perceptual states, which cause cognitive states. The latter are often assumed to consist (at least in part) of structured symbolic expressions. Cognitive computation over input symbol structures leads to output symbol structures that guide behavior. Behavior can have causal consequences for distal entities thus closing the causal loop.

Barsalou’s proposal is modest because it is concerned, essentially, with just one link in the chain, the link between perceptual states and symbol structures. Significantly, his proposal does not require a revised account of either the perceptual states that produce symbol structures or the cognitive computations that transform them. His core claim is that changing the nature of the symbol structures is sufficient to repair the deficiencies in amodal symbol systems theorizing. Thus he proposes simply that amodal symbol structures should be replaced with perceptual symbol structures which have two primary characteristics not shared with amodal structures. Perceptual symbol structures are hypothesized to be modality specific (or multi-modal) and analogical. The force of the latter point is that “the structure of a perceptual symbol corresponds, at least somewhat, to the perceptual state that produced it” (sect. 1.1). Thus the suggestion is that the distinction between perceptual symbols and amodal symbols is akin to the distinction between analogue and digital modes of representation. The primary requirements have a number of further consequences discussed in section 2.2. Perceptual symbols are dynamic not discrete, contextually variable, nonspecific and potentially indeterminate. They have, in other words, a degree of plasticity that is not characteristic of amodal symbols. They are, nevertheless, required to support the full range of computations needed for a fully fledged conceptual system. Most important in this regard, Barsalou argues that perceptual symbol structures must be capable of unbounded generativity and recursive elaboration (sect. 3.1).

Sympathetic though I am to the idea that perception and cognition should be better integrated, I think that Barsalou’s proposal is almost guaranteed not to work. It is essential for unbounded generativity and recursion that primitive symbol tokens be amodal and reliably identifiable independent of context. Barsalou proposes to give up these essential characteristics without a clear picture of how the much sloppier perceptual symbols are supposed to perform the same functions. He is right to emphasize the constraints on symbols arising from the neural foundations of the cognitive system, and right to note that the evidence for amodal symbols is not strong, but changing the nature of the symbols does not seem to be the solution to the problem. If a theory of cognitive processing requires unbounded generativity and recursion, then introducing sloppy symbols is a weakness. A better bet, if you believe that neural systems cannot support amodal symbols and the requisite compositional and recursive processes, is to question the nature of the requirement for productivity. It is also unclear how perceptual symbols solve the transduction problem. It is not clear, for example, that abstraction from perceptual states via selective attention is superior to any other current proposal in this regard.

Can one then characterize the requirement for productivity in a way that relieves the burden on internal representation and allows perceptual processes to play a larger part in cognitive processing generally? I believe that it can be done but it requires a more radical picture of the nature of cognitive architecture than

the one Barsalou paints. The traditional view of productivity relies primarily on internal resources for both storage and processing. Hence the requirement for internal structures that can store arbitrarily complex, iterated, nested symbolic expressions. It is quite plausible, however, given the existence of external symbol storage media such as books, that productivity might be understood as an interaction between external symbol structures and internal processes. In the most extreme case there need be no internal symbol storage. An argument in favor of models of cognitive architecture with external symbolic resources has been advanced by Wells (1998) based on a re-analysis of Turing's theory of computation. If one argues that cognitive architecture extends into the environment of the cognizer, cognition does indeed become integrally linked to perception. Moreover, the representational burden on internal states can be greatly reduced. Donald (1991; see also BBS multiple book review of Donald's *Origins of the Modern Mind* BBS 16(4) 1993) also emphasizes the significance of external symbol storage for theories of cognitive architecture.

Barsalou's proposal demonstrates the need for a theory of this kind. He, among others, has pointed out the substantial difficulties faced by theories based on amodal symbol systems. His achievement in the present target article is to have set out a comprehensive outline of what a theory based on perceptual symbols would be like. He has, so to speak, covered all the angles and indicated the central problems that any theory using perceptual symbols would have to solve. In doing this, he has, I think, shown that such a theory is unlikely to succeed.

ACKNOWLEDGMENT

I am grateful to Bradley Franks for discussion and advice.

Perceiving abstract concepts

Katja Wiemer-Hastings and Arthur C. Graesser

Department of Psychology, The University of Memphis, Memphis, TN 38152-6400. kwiemer@latte.memphis.edu graesser@cc.memphis.edu
www.psycc.memphis.edu/faculty/graesser/graesser.htm

Abstract: The meanings of abstract concepts depend on context. Perceptual symbol systems (PSS) provide a powerful framework for representing such context. Whereas a few expected difficulties for simulations are consistent with empirical findings, the theory does not clearly predict simulations of specific abstract concepts in a testable way and does not appear to distinguish abstract noun concepts (like *truth*) from their stem concepts (such as *true*).

Do perceptual symbol systems (PSS) solve the "challenge of any perceptual representation system," that is, can they account for the representation of abstract concepts? Abstract concepts have no perceivable referents. It is therefore challenging to specify their representation. It is well documented that comparatively few attributes are generated for abstract concepts in free generation tasks (Graesser & Clark 1985; Markman & Gentner 1993; McNamara & Sternberg 1983). Hampton (1981) has shown that many abstract concepts are not structured as prototypes.

With such difficulties in mind, it may seem somewhat surprising that a perceptually based theory should provide a compelling framework to represent abstract concepts. However, we argue that the nature of abstract concepts requires such an approach. Previous theories have emphasized that such concepts depend on context (e.g., Quine 1960). Grounding abstract concepts in context, however, only works if context provides useful information for their meaning representation. That is, context itself needs to have a representation that is grounded. And of course, the ultimate grounding mechanism of humans is perception.

Concept representations characterize a concept and distinguish it from other concepts. Barsalou demonstrates that PSS can in principle represent abstract concepts, but can it also handle their differentiation? This issue is important. Abstract concepts are of-

ten very similar to each other, therefore requiring very fine distinction. Examples for similar abstract concepts are *argument*, *disagreement*, and *conflict*. The similarity of such word sets should result in very similar simulations because the concepts occur in similar contexts. At the same time, however, we know that there are subtle differences in meaning, and such differences must be captured by representations. Do PSS capture such differences?

Suppose separate groups were asked to list attributes for either *argument* or for *disagreement*. Conceivably, both groups would come up with similar lists, suggesting identical representations. However, if asked to generate attributes for both concepts, they might attempt to list distinctive features via a careful comparison of simulations. This illustrates a problem: PSS have the power to represent abstract concepts, but the theory does not provide principled predictions for simulations of particular abstract concepts. Instead, ad hoc simulations account for a concept depending on a context or on contrast concepts.

Contexts of abstract concepts may vary considerably. Hampton (1981) pointed out that abstract concepts are "almost unlimited in the range of possible new instances that could be discovered or invented" (p. 153). This is especially true for a subset of abstract nouns that have a content, such as *belief*, *truth*, and *idea*. Content concepts reflect various degrees of abstraction, as illustrated for the concept of truth:

1. *I knew he was not telling me the truth.* (*Truth* is a concrete referent in a situation; it is a "fact.")
2. *I know that he always tells the truth.* (*Truth* is a feature that is shared by a set of statements.)
3. *Scientists want to know the truth.* (*Truth* is the general, absolute knowledge of how things are.)

Barsalou's suggested simulation for *truth* is appropriate for (1) and (2), in that it involves the comparison of facts and statements. For (3), however, it seems that the simulation needs to express a higher level of abstractness, because this *truth* (the set of facts) is as yet unknown. It is unclear how PSS would handle cases like (3).

A second interesting aspect of content concepts is that their content can vary infinitely. This may be associated with the observation that abstract concepts are difficult to define and overall take longer to process (for an overview see Schwanenflugel 1991). Consider the following expressions of ideas: *He might be in the other room, perhaps you need to dial a 1 first, let's have pizza*, and so on. How can an abstracted representation of *idea* be constructed from such examples? What do these ideas have in common?

They probably share contextual aspects, such as an initial problem: *I cannot find him, I cannot get through, or what should we eat?* Then, somebody expresses an idea. Consequently, different things may happen: An opposing reaction (*We had pizza yesterday*), a new behavior to solve the problem (*he was indeed in the other room, or dialing 1 does not help either*), or an alternative idea (*I've just looked there. Perhaps he is outside?*). Due to the variation, it may be difficult to identify which aspects are critically related to the word meaning and should therefore be used to frame the simulation. However, if concepts are presented in a specific context, the context can frame the simulation and provide clear referents. Compare the following contexts:

1. "It is too bad that he quit his job. He always had such good ideas!"
2. "I have lost my key somewhere."
 "Did you search your jacket pockets?"
 "No - great idea!"

Example (1) requires some abstraction of idea in the context of a job, whereas in (2), the idea is a concrete suggestion to do something. A simulation should be easier to construct in the second case because the situation sets up an event sequence. This does not present a problem, however, since it is consistent with the observation that abstract language is processed faster if presented in a meaningful context (Schwanenflugel & Shoben 1983).

Finally, it is unclear how a perceptual symbols system distinguishes between a stem concept (“true”) and a complex abstract noun (“truth”). It is true that the moon orbits the earth. Also, occasionally it is true that *we are tired*. We know that these statements are true by comparison with the facts. It seems that the simulations for *true* and *truth* would be the same. However, it could be argued that truth and something being true are very different concepts.

Perceptual symbols in language comprehension: Can an empirical case be made?

Rolf A. Zwaan, Robert A. Stanfield, and Carol J. Madden

Department of Psychology, Florida State University, Tallahassee, FL 32306-1270. freud.psy.fsu.edu:80/~zwaan/

Abstract: Perceptual symbol systems form a theoretically plausible alternative to amodal symbol systems. At this point it is unclear whether there is any truly diagnostic empirical evidence to decide between these systems. We outline some possible avenues of research in the domain of language comprehension that might yield such evidence. Language comprehension will be an important arena for tests of the two types of symbol systems.

Barsalou provides convincing theoretical arguments for why a perceptual symbol system should be preferred over amodal systems: greater parsimony, the ability to predict rather than postdict effects, resolution of the symbol-grounding problem, consistency with the notion of embodiment, and avoidance of an evolutionary quantum leap presupposition. He also discusses examples of empirical evidence. It is not clear to us at this point how diagnostic these are with respect to a comparison between perceptual and amodal symbol systems. For example, Barsalou (sect. 2.3) interprets neuroimaging evidence showing that visual areas of the brain are active during conceptual tasks involving animals and motor areas during conceptual tasks involving tools as support for perceptual symbol systems. Proponents of amodal systems might grant that representations are stored in the corresponding perceptual areas of the brain, but still maintain that these representations are amodal. The problem, as Barsalou (sect. 1.2.2) notes is that amodal systems might be able to explain (away) many findings on a post-hoc basis.

Because we find Barsalou’s view congenial and inspiring, we attempt to identify some existing and possible lines of research that might have diagnostic value with respect to the comparison between perceptual and amodal symbol systems and as such could bolster the case for perceptual symbol systems empirically. Amodal propositional systems are arguably more dominant in the area of language comprehension than in any other area of cognition. Therefore, it is a crucial arena for evaluating perceptual and amodal symbol systems.

We will first clarify our understanding of language comprehension in a perceptual-symbol framework, extrapolating from Barsalou’s comments. Language comprehension is tantamount to the construction of a mental model of the described state of affairs: a situation model (e.g., Zwaan & Radvansky 1998). Barsalou (sect. 2.4.2) argues that situation models are, in terms of a perceptual-symbol framework, equivalent to surface level simulations only.¹ We agree with this characterization. Situation models represent specific situations bound in time and space. As such, they are tokens, whereas related knowledge structures, such as scripts and frames, are types. We understand simulators to be types and simulations to be tokens. Thus, frames and scripts are assemblies of simulators (which differ in that the simulators become activated in a specified temporal order in scripts but not in frames), whereas situation models are simulations generated in part by frames and scripts. These simulations run more smoothly – and perhaps seem more perception-like – when their simulators can be assembled

and sequenced quickly. This happens when comprehenders have a great deal of background knowledge, for example in the form of frames and scripts. Thus, skilled text comprehension is analogous (but not identical!) to the perception of unfolding events. Such a perspective might be better equipped to explain how readers become immersed in narrative worlds (Zwaan, in press) than would amodal systems.

How could research on situation models yield evidence that is diagnostic with respect to a comparison between perceptual and amodal symbol systems? We sketch five existing and potential lines of research.

1. *Perceptual effects during comprehension.* Morrow and Clark (1989) demonstrated that readers’ interpretation of motion words, such as “approach” depends on perceptual characteristics of the situation. Furthermore, comprehenders provide faster responses to words denoting events and objects present in the situation compared to discontinued events and nonpresent objects (e.g., Macdonald & Just 1989; Zwaan 1996). There is no straightforward way to account for these findings using propositional analyses.

2. *Rapid integration of perceptual and language-based information.* Klatzky et al. (1989) found that the comprehension of verbally described actions is facilitated when subjects are allowed to first form their hand into a shape appropriate for the action (for example, fingers pinched together when reading about throwing a dart).

3. *Interference between secondary perceptual/imagery task and situation-model construction.* Fincher-Kiefer (1998) recently showed that comprehenders who had to keep a high-imagery sentence unrelated to the text in working memory had greater difficulty constructing a situation model than comprehenders who had to keep a low-imagery sentence in working memory.

4. *Effects of embodiment on comprehension.* We know of no research in this domain. Yet, this type of evidence might yield strong evidence against amodal systems. For example, it might be that a basketball player represents the situation described by “He picked up the basketball” differently (one-handed) from a smaller person (two-handed). One would also assume that women who have given birth construct different mental representations of a story about childbirth than other women or than men. Amodal systems would have no a priori way to account for this.

5. *Physiological effects of comprehension.* Clinical psychologists have amassed a great deal of relevant evidence, which is not discussed by Barsalou. For example, Lang and his colleagues have conceptualized and demonstrated how emotional imagery can have powerful physiological effects (Lang 1979).

Presumably, the evidence yielded by these lines of research would not be equivalent in diagnostic value. For example, the rapid integration of language-based and perceptual information and the physiological effects of comprehension are not necessarily predicted by amodal symbol systems, but might be accommodated by them post hoc, maybe with some additional assumptions. The occurrence of perceptual effects during comprehension would be more difficult to accommodate by amodal systems. Even more diagnostic would be demonstrations of interference of unrelated perceptual representations on the construction of situation models. Finally, the most diagnostic evidence, in our view, would be demonstrations of variable embodiment on comprehension processes. Perhaps none of these lines of research will yield convincing evidence for perceptual symbols independently, but taken together they might. At the very least, they would force proponents of amodal symbol systems to postulate an increasingly unwieldy and implausible range of special assumptions to postdict the findings.

Barsalou has provided a viable alternative against reigning amodal symbol systems. He has successfully integrated ideas and concepts from linguistics, philosophy, and neuroscience into a cognitive-psychological framework. The perceptual-symbol framework shines a searchlight on research topics like the five listed above, which have received scattered attention in the past. As such, it could provide a unified account of the extant and future

findings and thus influence the agenda for cognitive science in years to come.

NOTE

1. This might not be entirely accurate for mental models in general, because some are more generic and therefore more like simulators; see Johnson-Laird's (1983, pp. 422–29) typology.

Author's Response

Perceptions of perceptual symbols

Lawrence W. Barsalou

Department of Psychology, Emory University, Atlanta, GA 30322.

barsalou@emory.edu userwww.service.emory.edu/~barsalou/

Abstract: Various defenses of amodal symbol systems are addressed, including amodal symbols in sensory-motor areas, the causal theory of concepts, supramodal concepts, latent semantic analysis, and abstracted amodal symbols. Various aspects of perceptual symbol systems are clarified and developed, including perception, features, simulators, category structure, frames, analogy, introspection, situated action, and development. Particular attention is given to abstract concepts, language, and computational mechanisms.

I am grateful for the time and energy that the commentators put into reading and responding to the target article. They raised many significant issues and made many excellent suggestions. Even the strongest criticisms were useful in clarifying misunderstandings and in addressing matters of importance. I have organized the commentaries into two general groups: (1) those that defended amodal symbol systems and (2) those that addressed the properties of perceptual symbol systems. Because the second group was so large, I have divided it into four smaller sections. In the first, I address a wide variety of issues surrounding perceptual symbol systems that include perception, features, simulators, frames, introspection, and so forth. The final three sections address the topics raised most often: abstract concepts, language, and computational mechanisms.

R1. Defense of amodal symbol systems

No commentator responded systematically to the criticisms of amodal symbol systems listed in sections 1.2.2 and 1.2.3 of the target article. Little direct evidence exists for amodal symbols; it is difficult to reconcile amodal symbols with evidence from neuroscience; the transduction and grounding of amodal symbols remain unspecified; representing space and time computationally with amodal systems has not been successful; amodal symbol systems are neither parsimonious nor falsifiable. Expecting a single commentator to address all of these concerns in the space provided would certainly be unfair. I was struck, however, by how little attempt overall there was to address them. Having once been deeply invested in amodal symbol systems myself (e.g., Barsalou 1992), I understand how deeply these convictions run. Nevertheless, defending amodal symbol systems against such criticisms seems increasingly important to maintaining their viability.

A handful of commentators defended amodal systems in various other ways. Although most of these defenses were concrete and direct, others were vague and indirect. In addressing these defenses, I proceed from most to least concrete.

R1.1. Amodal symbols reside in sensory-motor areas of the brain. The defense that struck me as most compelling is the suggestion that amodal symbols reside in sensory-motor regions of the brain (Aydede; Zwaan et al.). A related proposal is that perceptual representations are not constitutive of concepts but only become epiphenomenally active during the processing of amodal symbols (Adams & Campbell). In the target article, I suggested that neuroscience evidence implicates sensory-motor regions of the brain in knowledge representation (sect. 2.1., 2.2, and 2.3). When a lesion exists in a particular sensory-motor region, categories that depend on it for the perceptual processing of their exemplars can exhibit knowledge deficits. Because the perception of *birds*, for example, depends heavily on visual object processing, damage to the visual system can produce a deficit in category knowledge. Neuroimaging studies of humans with intact brains similarly show that sensory-motor areas become active during the processing of relevant categories. Thus, visual regions become active when accessing knowledge of *birds*.

Several commentators noted correctly that these findings are consistent with the assumptions that (a) amodal symbols represent concepts and that (b) these symbols reside in sensory-motor regions (Adams & Campbell; Aydede; Zwaan et al.). If these assumptions are correct, then damage to a sensory-motor region could produce a deficit in category knowledge. Similarly, sensory-motor regions should become active when people with intact brains process categories.

This is an important hypothesis that requires careful empirical assessment. Damasio's (1989) convergence zones provide one way to frame the issue. According to his view, local associative areas in sensory-motor regions capture patterns of perceptual representation. Later, associative areas reactivate these sensory-motor representations to simulate experience and thereby support cognitive processing. As the quotation from Damasio (1989) in section 2.1 illustrates, he believes that reactivating sensory-motor representations is necessary for representing knowledge – activation in a nearby associative area never stands in for them. In principle, however, activation in local associative areas could stand in for sensory-motor representations during symbolic activity, thereby implementing something along the lines of amodal symbols, with perceptual representations ultimately being epiphenomenal.

Behavioral findings argue against this proposal. Studies cited throughout the target article show that perceptual variables predict subjects' performance on conceptual tasks. For example, Barsalou et al. (1999) report that occlusion affects feature listing, that size affects property verification, and that detailed perceptual form predicts property priming. The view that activation in associative areas represents concepts does not predict these effects or explain them readily. Instead, these effects are more consistent with the view that subjects are reactivating sensory-motor representations. Studies on language comprehension similarly exhibit such effects (sects. 4.1.6, R4.5, and Zwaan et al.).

Further empirical evidence is clearly required to resolve this issue with confidence. Neuroimaging studies that assess both sensory-motor and local associative areas during perception and conception should be highly diagnostic. Similarly, lesion studies that assess simultaneous deficits in perception and conception should be useful, as well. If sensory-motor representations are epiphenomenal, then patients should frequently be observed who exhibit deficits in sensory-motor representation but not in corresponding knowledge.

R1.2. The causal theory of concepts. Ignoring my explicit proposals to the contrary, **Aydede** claims that I addressed only the form of perceptual symbols and ignored their reference. From this flawed premise, he proceeds to the conclusion that a causal theory of amodal symbols must be correct. On the contrary, the target article clearly addressed the reference of perceptual symbols and agreed with Aydede that causal relations – not resemblances – are critical to establishing reference. Section 2.4.3 states that simulators become causally linked to physical categories, citing Millikan (1989). Similarly, section 3.2.8 is entirely about the importance of causal mechanisms in establishing the reference of perceptual symbols, citing Dretske (1995). Indeed, this section makes Aydede's argument that resemblance is inadequate for establishing a symbol's reference. Although resemblance is often associated with perceptual views of knowledge, I took great pains to dissociate the two.

Although I would have valued **Aydede's** analysis of my actual proposal, reminding him of it provides an opportunity for elaboration. I agree with the many philosophers who hold that much conceptual knowledge is under the causal or nomic control of physical categories in the environment. Indeed, this assumption underlies all learning theories in cognitive science, from exemplar theories to prototype theories to connectionism. In this spirit, section 2.4.3 proposed that simulators come under the control of the categories in the world they represent. As a person perceives the members of a category, sensory-motor systems are *causally* driven into particular states, which are captured by various associative areas. The target article does not specify how these causal processes occur, disclaiming explicitly any attempt to provide a theory of perception (sect. 2, para. 1–3). Nor does it make any claim that perceptual states resemble the environment (although they do in those brain regions that are topographically mapped). It is nevertheless safe to assume that causal processes link the environment to perceptual states. The critical claim of perceptual symbol systems is that these causally produced perceptual states, whatever they happen to be, constitute the representational elements of knowledge. Most critically, if the environment is causally related to perceptual states, it is also causally related to symbolic states. Thus, perceptual symbol systems constitute a causal theory of concepts. Simulators arise in memory through the causal process of perceiving physical categories.

In contrast, consider the causal relations in the amodal theories that **Aydede** apparently favors (as do **Adams & Campbell**). According to these views, physical entities and events become causally related to amodal symbols in the brain. As many critics have noted, however, this approach suffers the problems of symbol grounding. Exactly what is the process by which a physical referent becomes linked to an amodal symbol? If philosophers want to take causal mech-

anisms seriously, they should provide a detailed process model of this causal history. In attempting to do so, they will discover that perception of the environment is essential (Harnad 1987; 1990; Höffding 1891; Neisser 1967), and that some sort of perceptual representation is necessary to mediate between the environment and amodal symbols. Aydede notes that perceptual representations could indeed be a part of this sequence, yet he fails to consider the implication that potentially follows: Once perceptual representations are included in the causal sequence, are they sufficient to support a fully functional conceptual system? As I argue in the target article, they are. If so, then why include an additional layer of amodal symbols in the causal sequence, especially given all the problems they face? Nothing in Aydede's commentary makes a case for their existence or necessity.

R1.3. Supramodal systems for space and time. Two commentaries suggest that spatial processing and temporal processing reside in amodal systems (**Freksa et al.; Mitchell & Clement**). According to this view, a single spatial system subserves multiple modalities, allowing coordination and matching across them. Thus, the localization of a sound can be linked to an eye movement that attempts to see the sound's source. Similarly, a single temporal system subserves multiple modalities, coordinating the timing of perception and action. Because the spatial and temporal systems are not bound to any one modality, they are amodal, thereby constituting evidence for amodal symbol systems.

To my knowledge, the nature of spatial and temporal systems remains an open question. To date, there is no definitive answer as to whether each modality has its own spatial and temporal processing system, or whether general systems serve all modalities. Indeed, *both* modality-specific and modality-general systems may well exist. Ultimately, this is an empirical question.

Imagine, however, that domain-general systems *are* found in the brain. In the framework of the target article, they would *not* constitute amodal systems. As defined in section 1.2, an amodal symbol system contains representations that are transduced from perceptual states – they do not contain representations from perception. If domain-general systems for perceiving space and time exist, their representations do not satisfy this definition for amodal symbols. To see this, consider the representations that would arise in domain-general systems during perception. Although these representations are not modality-specific, they are nevertheless perceptual. They constitute fundamental parts of perception that integrate the specific modalities – what I will call *supramodal representations*. To appreciate why these representations are not amodal, consider the representation of space and time during conceptual processing. If supramodal representations in perception are also used to represent space and time during conceptual processing, they constitute perceptual symbols. Amodal symbols for space and time would only exist if supramodal representations during perception were transduced into a new representation language that supports conceptual processing. Should no such transduction process occur, supramodal representations of time and space form central components of perceptual symbol systems.

R1.4. Latent semantic analysis (LSA). Based on my arguments about artificial intelligence (sect. 4.4), **Landauer**

takes me to mean that no silicon-based system could acquire human knowledge. He then shows how computers that implement LSA mimic humans on a wide variety of knowledge-based tasks. I never claimed, however, that amodal systems cannot mimic humans behaviorally! Indeed, section 1.2 noted that amodal symbols have so much power that they can probably mimic all human behavior. My claim in section 4.4 was different, namely, if knowledge is grounded in sensory-motor mechanisms, then humans and computers will represent knowledge differently, because their sensory-motor systems differ so radically. I acknowledged the possibility that future technological developments might produce computers with sensory-motor systems closer to ours, in which case their knowledge could be represented similarly. No one is currently under the illusion, however, that the input-output systems of today's computers are anything like human sensory-motor systems. Thus, if knowledge is implemented in perceptual systems, it would have to take very different forms.

In making his case for LSA, **Landauer** fails to address the problems I raise for amodal symbol systems in section 1.2: (1) What is the direct evidence that co-occurrence frequency between words controls performance? A system based on co-occurrence frequencies can mimic human behavior, thereby providing indirect evidence that this construct is useful, but is there evidence implicating this basic unit of analysis directly? (2) What is the neural story? Do lesion and neuroimaging data provide support? (3) How are perception and conception linked? Everyone would agree that a conceptual system helps us interpret perceived events in the world, but how do word co-occurrences map onto perceived individuals in the world? When we perceive an object, how is it categorized using LSA? When we conceptualize an object, how do we find its referent in the world? (4) How does LSA represent knowledge about space and time? It is difficult to see how a system that only tracks word co-occurrence can represent these basic aspects of perception. Other concerns about LSA arise, as well. How does LSA accomplish the conceptual functions addressed in the target article, such as distinguishing types and tokens, implementing productivity, and representing propositions? Using only word correlations, it does not seem likely that it can implement these basic human abilities. Finally, are co-occurrences tracked only between words, or between corresponding amodal symbols, as well? If the latter is the case, what is the nature of this amodal system, and how is it related to language?

It is important to note that correlation – not causation – underlies LSA's accounts of human data. Because word co-occurrence exists in the world and is not manipulated experimentally, infinitely many potential variables are confounded with it. To appreciate this problem, recall how LSA works. On a given task, the likelihood of a particular response (typically a word) is predicted by how well it correlates with the overall configuration of stimulus elements (typically other words). Thus, responses that are highly correlated with stimulus elements are more likely to be produced than less correlated responses. The problem is that many, many other potential variables might well be correlated with word co-occurrence, most notably, the perceptual similarity of the things to which the words refer. It is an interesting finding that computer analyses of word correlations produce a system that mimics human behaviors. It does not follow, though, that these correlations are causally

important. Of course the human brain could work like LSA, but to my knowledge, no evidence exists to demonstrate this, or to rule out that variables correlated with word co-occurrence are the critical causal factors.

A number of empirical findings suggest that knowledge is not grounded in word co-occurrence. First, aphasics typically lose only language; they do not lose knowledge (**Lowenthal**). If knowledge were simply represented in word co-occurrence, they should lose knowledge, as well. Second, **Glenberg** reports that perceptual affordances enter into language comprehension, and that word co-occurrence cannot explain these results (also see sects. 4.1.6 and R4.5). **Landauer** believes that this is because words for certain idioms, metaphors, and modern phrases have not been included in the language corpus that underlies LSA, but again the problem of correlation versus causation arises. Landauer wants to believe that these missing entries have causal significance, when actually their status is merely correlational. Because Glenberg's affordances are also correlational, more powerful laboratory techniques are needed to tease these factors apart. To the extent that sensory-motor affordances are implicated in conceptual processing, however, LSA has no way of handling them. In this spirit, Solomon and Barsalou (1999a; 1999b) controlled the strength of lexical associations in the property verification task and found effects of perceptual variables. In a task involving only linguistic stimuli, nonlinguistic factors significantly affected performance when linguistic factors were held constant.

One way to extend LSA would be to apply its basic associative mechanism to perceptual elements. Just as co-occurrences of words are tracked, so are co-occurrences in perception. Depending on the specific details of the formulation, such an approach could be virtually identical to the target article's proposal about frames. As described in section 2.5., a frame accumulates components of perception isolated by selective attention, with the associative strength between components reflecting how often they are processed together. Thus, the co-occurrence of perceptual components lies at the heart of these perceptually grounded frames. Notably, however, it is proposed that these structures are organized by space and time, not simply by associative strength. To later simulate perceptual experience, the mere storage of associations will not be sufficient – extensive perceptual structure will be necessary, as well. If so, a simple extension of LSA from words to perception probably will not work.

Furthermore, we have known since Chomsky (1957) that the human brain is not simply a generalization mechanism. Besides generalizing around known exemplars, the brain produces infinite structures from finite elements productively. Not only is this true of language, it is also true of conception (Fodor & Pylyshyn 1988). People can conceive of concepts that they have never encountered and that are not generalizations of past concepts. Neither a word-based nor a perception-based version of LSA can implement productivity. There is no doubt that the human brain is an associative device to a considerable extent, yet it is also a productive device. Although LSA may explain the associative aspects of cognition, it appears to have no potential for explaining the productive aspects.

Finally, in a superb review of picture-word processing, Glaser (1992) concluded that a mixture of two systems accounts for the general findings in these paradigms. Under

certain conditions, a system of word associations explains performance; under other conditions, a conceptual system with a perceptual character explains performance; under still other conditions, mixtures of these two systems are responsible. Based on my own research, and also on reviews of other relevant literature, I have become increasingly convinced of Glaser's general thesis, although perhaps in a somewhat more radical form: People frequently use word associations and perceptual simulations to perform a wide variety of tasks, with the mixture depending on task circumstances (e.g., Barsalou et al. 1999; Solomon & Barsalou 1999a; 1999b; Wu & Barsalou 1999). Furthermore, I have become convinced that LSA provides an excellent account of the word association system (as does Hyperspace Analogue to Language, HAL; Burgess & Lund 1997). When people actually use word associations to perform a task, LSA and HAL provide excellent accounts of the underlying processing. However, when a task demands true conceptual processing, or when the word association strategy is blocked, the conceptual representations required are implemented through perceptual simulation. Although such simulations may rely on structures of associated perceptual components, they also rely on nonassociative mechanisms of perception and productivity.

R1.5. Amodal symbols arise through abstraction. Adopting the framework of connectionism, **Gabora** argues that amodal symbols arise as similar experiences are superimposed onto a common set of units. As idiosyncratic features of individual experiences cancel each other out, shared features across experiences coalesce into an attractor that represents the set abstractly. I strongly agree that such learning mechanisms are likely to exist in the brain, and that they are of considerable relevance for perceptual symbol systems. Most important, however, this process does not satisfy the definition of amodal symbols in section 1.2. To the extent that Gabora's learning process occurs in the original units that represent perception, it implements a perceptual symbol system, even if generic attractors evolve over learning. As described in section 2.5 and shown in Figure 3, repetitions of the same perceptual symbol are superimposed to produce increasingly generic components of frames. Conversely, Gabora's learning process would only become an amodal symbol system if a population of nonperceptual units were to recode perceptual activity, as in a feedforward net with hidden units (sect. 1.2). As long as no such recoding takes place, however, I view Gabora's proposal as highly compatible with the proposal that abstractions are grounded perceptually. **Schwartz et al.** make virtually the same proposal as Gabora but view it as an implementation of a perceptual symbol system, not as an implementation of an amodal symbol system.

R1.6. General skepticism. The commentaries by **Landau** and **Ohlsson** raise specific criticisms of perceptual symbol systems that I will address later. Here I address their general belief that amodal symbol systems must be correct. Landau and Ohlsson both view perceptual symbol systems as essentially a classic empiricist theory. Because they believe that prior empiricist theories have been completely discredited, they believe that perceptual symbol systems could not possibly work either.

First, these commentaries prove the point I made in the target article that perceptual approaches to knowledge are

often oversimplified for the sake of criticism (sect. 1.3). I do not have the sense that **Landau** and **Ohlsson** have appreciated this point, or have made much of an attempt to entertain anything but a simplistic perceptual view. I continue to believe that when more substantial and ambitious perceptual theories are entertained, they are not only plausible and competitive, they are superior.

Second, the theory in the target article departs significantly from classic empiricist theories. Whereas prior theories focused on consciously experienced images, perceptual symbol systems focus on neural states in sensory-motor systems (sect. 2.1). To my mind, this is the most exciting feature of this proposal, primarily because it leads one to think about the perceptual view in provocative new ways.

Third, I explicitly stated that my theory is not a simple empiricist view of knowledge, appealing to genetically regulated mechanisms that structure brain development (sect. 2.4.1). In particular, I proposed that genetically based mechanisms for space, objects, movement, and emotion organize perceptual symbols. **Landau** and **Ohlsson** seem to believe that nativism is the answer to the ills of empiricism. However, it is one thing to view nativism naively as a panacea for every difficult problem, and quite another to specify the complex epigenetic processes that characterize the genetic regulation of development (Elman et al. 1996).

Fourth, perceptual symbol systems go considerably further than previous theories in laying out a fully functional theory of concepts, largely because the criteria for such a theory were not known until modern times. To my knowledge, no prior perceptual theory of knowledge has accomplished this. Although prior theories have included mechanisms that had potential for implementing these functions (sect. 1.3), they were typically not developed in these ways.

Finally, in embracing amodal symbol systems as the only reasonable alternative, **Landau** and **Ohlsson** fail to address the problems raised for amodal symbol systems in section 1.2. Without resolving these major problems, it is difficult to see how anyone could have such complete confidence in a theory.

R2. Specifying perceptual symbol systems

The majority of the commentaries addressed specific aspects of perceptual symbol systems. Some commentaries attempted to show that empirical findings or theoretical arguments support some aspect of the theory. Others developed a particular aspect in a positive manner. Still other commentaries attempted to undermine the theory by raising problems for some aspect. I have organized these remarks around each aspect of the theory addressed. For each, I have combined the relevant commentaries with the hope of clarifying and developing that part of the theory.

R2.1. Perception, perceptual symbols, and simulation.

Quite remarkably, **Aydede** claims that there is no way to distinguish perceptual and nonperceptual systems independently of modal and amodal symbols. Given the huge behavioral and neural literatures on perception, this is a rather surprising claim. A tremendous amount is known about the perceptual systems of the brain independent of anything anyone could possibly say about conceptual systems. Clearly, we know a lot about perceptual abilities and the neural mechanisms that implement them. Identifying these

abilities and mechanisms independently of perceptual symbols is not only possible but has long since been accomplished in the scientific disciplines that study perception. In defining perceptual symbols, I simply used these well-recognized findings. Thus, section 2.1 proposed that well-established sensory-motor mechanisms not only represent edges, vertices, colors, and movements in perception but also in conception. Perceptual symbols follow the well-trodden paths of perception researchers in claiming that the representational mechanisms of perception are the representational mechanisms of knowledge. What is perceptual about perceptual symbol systems is perception in the classic and well-established sense. The exceptions are introspective symbols (sect. 2.3.1), whose neural underpinnings remain largely unexplored (except for emotion and motivation).

Brewer similarly worries that perceptual symbols, neurally defined, are of little use because we currently have no knowledge of the neural circuits that underlie them. Again, however, we do have considerable knowledge of the neural circuits that underlie sensory-motor processing. Because perceptual symbols are supposed to use these same circuits to some extent, we *do* know something about the neural configurations that purportedly underlie them.

Although perceptual symbols use perceptual mechanisms heavily, this does not mean that conceptual and perceptual processing are identical. In the spirit of **Lowenthal's** commentary, important differences must certainly exist. As noted in sections 2.1 and 4.3 of the target article, perceptual symbols probably rely more heavily on memory mechanisms than does standard perception. Furthermore, the thesis of section 2.4.7 is that a family of representational processes capitalizes on a common set of representational mechanism in sensory-motor regions, while varying in the other systems that use them. This assumption is also implicit throughout the subsections of section 4.1, which show how basic cognitive abilities could draw on the same set of representational mechanisms. The point is not that these abilities are identical, only that they share a common representational basis.

Mitchell & Clement claim that I never define "simulation" adequately. It is clear throughout the target article, however, that a simulation is the top-down activation of sensory-motor areas to reenact perceptual experience. As section 2.4 describes in some detail, associative areas reactivate sensory-motor representations to implement simulations.

Wells claims that the extraction of perceptual symbols (sect. 2.2) is not an improvement over the unspecified transduction process associated with amodal symbols (sect. 1.2). The improvement, as I see it, is that the extraction of perceptual symbols is a well-specified process. It is easy to see how selective attention could isolate and store a component of perception, which later functions as a perceptual symbol. In contrast, we have no well-specified account of how an amodal symbol develops.

Finally, I would like to reiterate one further point about my use of "perception." The use of this term in "perceptual symbol systems" is *not* meant to imply that movements and action are irrelevant to knowledge. On the contrary, many examples of perceptual symbols based on movements and actions were provided throughout the target article. Rather, "perception" refers to the perception of experience from which perceptual symbols are extracted, including the perception of movement and action. Thus, "perception" refers to the *capture* of symbolic content, not to the *content* itself.

For additional discussion of perceptual representations in knowledge, see Brewer and Pani (1983) and Pani (1996).

R2.2. Features and attention. One could take **Wells's** concern about perceptual symbol extraction being unconstrained to mean that it provides no leverage on the problem of what can be a feature. **Landau** makes this criticism explicitly. Notably, the amodal account that Landau apparently prefers places no constraints on features at all. As the introduction to section 2 in the target article noted, amodal theories have not resolved the feature issue themselves. It is puzzling that what counts as evidence against the perceptual view is not also counted as evidence against the amodal view.

Furthermore, as **Landau** notes, perceptual views contain the seeds of a solution to this problem. Specifically, she notes that classic empirical theories define conceptual features as features that arise from the interaction of the sense organs with the physics of the environment. In other words, the features of conception are the features of perception. Clearly, not all sensory distinctions make their way to the conceptual level, for example, opponent color processing and center-surround inhibition. Yet many conceptual features do appear to have their roots in sensory systems, including the features for shape, color, size, and texture that occur frequently in reports of conceptual content. Although Landau does not credit perceptual symbol systems with drawing on perception for features, they certainly do, given the explicit mention of perceptually derived features in section 2.1.

Because of its complexity, I did not attempt to resolve the feature problem in the target article. The features of conception clearly go beyond perceptual features, as I noted in Note 5. Most important, high-level goals and conceptual structures capture aspects of perception that become features through selective attention. In most cases, the information selected is not a feature computed automatically in sensory processing but a higher-order configuration of basic sensory-motor units. It was for this reason that I cited Schyns et al. (1988) as showing that conceptual features clearly go beyond perceptual features.

The architecture of perceptual symbol systems is ideally suited to creating new features that exceed sensory-motor features. Because high-level frames and intuitive theories are represented as perceptual simulations (sect. 2.5 and 4.1.2), they are well suited for interacting with perception to isolate critical features. Furthermore, the symbol formation process in sect. 2.2 is well suited for creating new features through selective attention. A great deal clearly remains to be specified about how this all works. It is much less clear how amodal symbol systems accomplish the construction of novel features, again because they lack an interface between perception and cognition.

Schwartz et al. argue that I fail to specify how attention is controlled during the formation of perceptual symbols, such that perceptual symbols are as arbitrary as amodal symbols. Again, however, perceptual systems provide natural constraints on the formation of perceptual symbols. One prediction is that when empirical studies examine this process, they will find that perceptual symbols exhibit a weak bias for following the structure of perception. By no means, however, must perceptual symbols follow only this structure. Because selective attention is so flexible, it can focus on complex configurations of perceptual information

to establish symbols that serve higher goals of the system. The open-endedness of the symbols that people use should not be viewed as a weakness of perceptual symbol systems. Instead, it is something that any theory needs to explain, and something that perceptual symbol systems explain through the flexible use of selection attention (sect. 2.2). It would be much more problematic if the theory were simply limited to symbols that reflect the computations of low-level sensory-motor processing.

Ultimately, **Schwartz et al.** have another agenda. They argue that schematic category representations result from recruitment and fractionation, not from selective attention. It is implicit in their view that holistic perceptual states are recorded during learning, with differences across them canceling out to produce schematic representations of perceptual symbols. Similarly, **Gabora** argues for the importance of these mechanisms. Although I endorse these mechanisms, I strongly disagree with the view that selective attention plays no role in learning. As discussed in sections 2.2 and 4.1.3, selective attention has been strongly implicated in learning across decades of research in many areas. Furthermore, it is naive to believe that holistic recordings of perception occur, as **Hochberg** argues eloquently (also see Hochberg 1998). Being highly flexible does not make selective attention irrelevant to conceptual processing. Instead, it is absolutely essential for creating a conceptual system (sects. 1.4 and 4.1.3).

Hochberg and **Toomela** each propose that active information-seeking and goal pursuit guide selective attention and hence the acquisition of perceptual symbols. **Glenberg's** arguments about the importance of situated action are also consistent with this view. I agree strongly that top-down goal-achievement often guides selective attention, such that perceptual information relevant to these pursuits becomes acquired as perceptual symbols. However, bottom-up mechanisms may produce perceptual symbols, as well. For example, attention-grabbing onsets and changes in perception may attract processing and produce a perceptual symbol of the relevant content. Similarly, the perceptual layout of a physical stimulus may guide attention and the subsequent extraction of perceptual symbols. In this spirit, **Brewer** notes that physical proximity better predicts recall order for the objects in a room than does their conceptual relatedness. Thus, bottom-up as well as top-down factors are probably important to selecting the content of perceptual symbols. Furthermore, appealing to the importance of top-down factors pushes the difficult issues back a level. Attention is clearly under the control of goals and action to a considerable extent, yet how do we characterize the mechanisms underlying goals and actions that guide attentional selection?

R2.3. Abstraction, concepts, and simulators. I am the first to admit that specifying the mechanisms underlying simulators constitutes perhaps the most central challenge of the theory, and I noted this throughout the target article. Later I will have more to say about implementing simulators in specific mechanisms (sect. R5.3). Here I focus on some misunderstandings about them.

Without citing any particular statement or section of the target article, **Ohlsson** claims that I equated selection with abstraction. According to him, I believe that selectively storing a wheel while perceiving a car amounts to creating an abstraction of cars in general. Nowhere does the target

article state that selection is abstraction. On the contrary, this is a rather substantial and surprising distortion of my view. What I did equate roughly with selection was schematization, namely, the information surrounding focal content is filtered out, leaving a schematic representation of the component (sect. 2.2). However, I never claimed that a schematization is an abstraction, or that it constitutes a concept. Rather, I argued that the integration of many schematic memories into a simulator is what constitutes a concept (sect. 2.4). Actually, I never said much about abstraction per se but focused instead on the development of types (sects. 2.4.3. and 3.2.1). If I were to characterize abstraction in terms of perceptual symbol systems, I would define it as the development of a simulator that reenacts the wide variety of forms a kind of thing takes in experience. In other words, a simulator contains broad knowledge about a kind of thing that goes beyond any particular instance.

Adams & Campbell claim that I characterize simulators in terms of the functions they perform, such as categorization and productivity. Most basically, however, I defined simulators as bodies of knowledge that generate top-down reenactments or simulations of what a category is like (sect. 2.4). Clearly, this is a functional specification, but it is not the one that Adams & Campbell attribute to me. Adams & Campbell further complain that I fail to say anything about the mechanisms that underlie simulators (so do **Dennett & Viger** and **Schwartz et al.**). In section 2.4, however, I defined a simulator as a frame plus the mechanisms that operate on it to produce specific simulations. In section 2.5, I provided quite a bit of detail on the frames that underlie simulators. I acknowledged that this account is far from complete, and I agree that developing a full-blown computational theory is extremely important. For the record, though, I certainly did say something about the mechanisms that underlie simulators, and I continue to believe that these ideas provide useful guidance in developing specific accounts (sect. R5.1).

Landau believes that simulations are not capable of capturing the underlying content or abstractions of concepts. If I understand correctly, she is making the same argument **Adams & Campbell** made earlier that perceptual knowledge is epiphenomenal and not constitutive of concepts. In a sense, this is also similar to **Ohlsson's** argument that storing part of a category instance does not represent the respective category. Again, however, the claim is that an *entire simulator* – not a specific perceptual representation – comes to stand for a category. As just discussed, a simulator captures a wide variety of knowledge about the category, making it general, not specific. Furthermore, the abstraction process described by **Gabora** and by **Schwartz et al.** enters into the frames that underlie simulators, thereby making them generic (sects. R1.5 and R2.4). Finally, a simulator stands in a causal relation to its physical category in the world (sect. R1.2). On perceiving an instance of category, a simulator is causally activated, bringing broad category knowledge to bear on the instance (sects. 2.4.3 and 3.2.1). For all these reasons, simulators function as concepts.

Finally, **Siebel** takes issue with my claim that bringing the same simulator to bear across different simulations of a category provides stability among them (sect. 2.4.5). Thus, I claimed that the different simulations a person constructs of birds are unified because the same simulator for *bird* produced all of them. As Siebel correctly notes, each time

a simulator becomes active, the perceptual symbols processed becomes stored or strengthened, thereby changing the simulator's content. Thus, the same simulator – in terms of its content – cannot be brought to bear on different simulations to provide stability. “Same,” however, does not refer to content but to the body of knowledge that stands in causal relation historically to a physical category. Stability results from the fact that one particular body of knowledge typically becomes active when a particular category is conceptualized, even though that body of knowledge changes in content over time. Because all conceptualizations can be linked to this particular body of knowledge historically, they gain stability.

R2.4. Conceptual essences and family resemblances.

Lurking behind the concern that simulators cannot represent concepts may well be the belief that concepts have essences. When **Ohlsson, Landau, and Adams & Campbell** question whether perceptual representations capture underlying constitutive content, I take them to be worrying about this. Since Wittgenstein (1953), however, there have been deep reservations about whether concepts have necessary and sufficient features. For this reason, modern theories often portray essences as people's *naive belief* that necessary and sufficient features define concepts, even when they actually do not (e.g., Gleman & Diesendruck, in press). Should a concept turn out to have defining features, though, a simulator could readily capture them. If a feature occurs across all the instances of a category, and if selective attention always extracts it, the feature should become well established in the frame that underlies the simulator (sect. 2.5, Fig. 3; sect. R1.5; **Gabora; Schwartz et al.**). As a result, the feature is almost certain to become active later in simulations of the category constructed. Simulators have the ability to extract features that are common across the instances of a category, should they exist.

Research on the content of categories indicates that few features if any are typically common to all members of a category (e.g., Malt 1994; Malt et al., in press; Rosch & Mervis 1975). However, a small subset of features is likely to be true of many category members, namely, features that are characteristic of the category but not defining. Should a family resemblance structure of this sort exist in a category, perceptual symbol systems are quite capable of extracting it. Again, following the discussion of frames (sect. 2.5, Fig. 3) and abstraction (sects. R1.5 and R2.3), the more often a feature is extracted during the perceptual processing of a category, the better established it becomes in the category's frame. To the extent that certain features constitute a statistical regularity, the frame for the category will capture this structure and manifest it across simulations constructed later. The more a feature is processed perceptually, the more it should occur in category simulations.

Finally, **Ohlsson** seems to believe that a perceptual view of concepts could only work if common perceptual features underlie all instances of a concept. Nowhere did I claim this, and there is no reason it must be true. As just described, the frame that underlies a simulator captures the statistical distribution of features for a category, regardless of whether it possesses common features. For a category lacking common features, its simulator will produce simulations that do not share features with all other simulations but are related instead by a family resemblance structure.

Fauconnier raises the related point that a word may not

be associated with a single simulator, taking me to mean that a simulator only produces simulations that share common features. Again, however, there is no a priori reason that a simulator cannot produce disjunctive simulations. Because different perceptions may be associated with the same word, different perceptual symbols may be stored disjunctively in the frame of the associated simulator and may later produce disjunctive simulations. As Fauconnier suggests, the content of a simulator is accessed selectively to project only the information relevant in a particular context, with considerably different information capable of being projected selectively from occasion to occasion (sect. 2.4.3).

R2.5. Frames and productivity. Barsalou and Hale (1993) propose that all modern representation schemes evolved from one of two origins: propositional logic or predicate calculus. Consider some of the differences between these two formalisms. Most basically, propositional logic contains binary variables that can be combined with simple connectives, such as *and*, *or*, and *implies*, and with simple operators, such as *not*. One problem with propositional logic is its inability to represent fundamentally important aspects of human thought and knowledge, such as conceptual relations, recursion, and bindings between arguments and values. Predicate calculus remedied these problems through additional expressive power.

Barsalou and Hale show systematically how many modern representation schemes evolved from propositional logic. By allowing binary variables to take continuous forms, fuzzy logic developed. By replacing truth preservation with simple statistical relations between variables, the construct of a feature list developed, leading to prototype and exemplar models. By adding complicated activation and learning algorithms to prototype models, connectionism followed. Notably, neural nets embody the spirit of propositional logic because they implement simple continuous units under a binary interpretation, linked by simple connectives that represent co-occurrence. In contrast, consider representation schemes that evolved from predicate calculus. Classic frame and schema theories maintain conceptual relations, binding, and recursion, while relaxing the requirement of truth preservation. Perceptual symbol systems similarly implement these functions (sect. 2.5.1), while relaxing the requirement that representations be amodal or language-like.

It is almost universally accepted now that representation schemes lacking conceptual relations, binding, and recursion are inadequate. Ever since Fodor and Pylyshyn's (1988) classic statement, connectionists and dynamic systems theorists have been trying to find ways to implement these functions in descendants of propositional logic. Indeed, the commentaries by **Edelman & Breen, Markman & Dietrich**, and **Wells** all essentially acknowledge this point. The disagreement lies in how to implement these functions.

Early connectionist formulations implemented classic predicate calculus functions by superimposing vectors for symbolic elements in predicate calculus expressions (e.g., Pollack 1990; Smolensky 1990; van Gelder 1990). The psychological validity of these particular approaches, however, has never been compelling, striking many as arbitrary technical attempts to introduce predicate calculus functions into connectionist nets. More plausible cognitive accounts that rely heavily on temporal asynchrony have been sug-

gested by Shastri and Ajjanagadde (1993) and Hummel and Holyoak (1997). My primary concern with this approach is that it is amodal and therefore suffers from the problems noted in section 1.2.

Edelman & Breen suggest that spatial relations provide considerable potential for implementing conceptual relations, binding, and recursion. Furthermore, they note that these functions achieve productivity through the combinatorial and recursive construction of simulations. As sections 2.5 and 3.1 indicate, I agree. Perceptual symbol systems provide more than a means of representing individual concepts; they provide powerful means of representing relations between them. My one addition to Edelman & Breen's proposal would be to stress the importance of temporal relations – not just spatial relations – as the basis of frames and productivity. The discussion of ad hoc categories in section 3.4.4 illustrates this point. One can think of a simulation as having both spatial and temporal dimensions. For example, a simulation of standing on a chair to change a light bulb contains entities and events distributed over both space and time. As suggested in section 3.4.4, ad hoc categories result from disjunctively specializing space-time regions of such simulations. Thus, the entity that is stood on to change a light bulb constitutes a space-time region, which can be specialized with simulations of different objects. Similarly, the agent standing on the chair can be specialized differently, as can the burned-out light bulb, the new light bulb, the fixture, and so forth. Each region that can be specialized constitutes a potential attribute in the frame for this type of event. Note that such regions are not just defined spatially, they are also defined temporally. Thus, the regions for the burned out bulb, the new bulb, and the chair all occupy different regions of time in the simulation, not just different regions of space. Because there are essentially an infinite number of space-time regions in a simulation, there are potentially an infinite number of frame attributes (Barsalou 1992; 1993).

Markman & Dietrich observe that frame structure arises in perceptual symbol systems through nativist mechanisms. With appropriate acknowledgment of epigenesis (Elman et al. 1996) and eschewing genetic determinism, I agree. Frame structure arises from the basic mechanisms of perception. Indeed, I have often speculated informally that predicate calculus also originated in perception. Markman & Dietrich further speculate that language is not the origin of frame structure. Conversely, one might speculate that the frame-like structure of language originated in perception as well. Because perception has relations, binding, and recursion, language evolved to express this structure through verbs, sentential roles, and embedded clause structure, respectively. Finally, Markman & Dietrich suggest that frame structure is not learned, contrary to how certain learning theorists might construe its origins in neural nets. I agree that the basic potential for frames lies deep in the perceptual architecture of the brain. No doubt genetic regulation plays an important role in producing the attentional, spatial, and temporal mechanisms that ultimately make frames possible. Yet, I hasten to add that specific frames are most certainly learned, and that they may well obey connectionist learning assumptions during their formation. Again, though, these connectionist structures are not amodal but are instead complex spatio-temporal organizations of sensory-motor representations.

Finally, **Wells** agrees that frame structure is important.

Without providing any justification, though, he claims that only amodal symbol systems can implement productivity, not perceptual symbol systems. Perhaps perceptual symbol systems do not exhibit exactly the same form of productivity as amodal symbol systems, but to say that they do not implement any at all begs the question. As section 3.1 illustrated in considerable detail, the hierarchical composition of simulations implements the combinatoric and recursive properties of productivity quite clearly. Wells then proceeds to argue that productivity does not reside in the brain but resides instead in interactions with the environment. This reminds me of the classic behaviorist move to put memory in the environment as learning history, and I predict that Wells's move will be as successful.

I sympathize with **Wells's** view that classic representational schemes have serious flaws. As **Hurford** notes, however, it is perhaps unwise to throw out the symbolic baby with the amodal bath water. There are many good reasons for believing that the brain is essentially a representational device (Dietrich & Markman, in press; Prinz & Barsalou, in press b). Some theorists have gone so far as arguing that evolution selected representational abilities in humans to increase their fitness (Donald 1991; 1993). Perceptual symbol systems attempt to maintain what is important about representation while similarly attempting to maintain what is important about connectionism and embodied cognition (sect. R6). The brain is a statistical, embodied, *and* representational device. It is our ability to represent situations offline, and to represent situations contrary to perception, that make us such amazing creatures.

R2.6. Analogy, metaphor, and complex simulations. As noted by **Markman & Dietrich**, the ability of perceptual symbol systems to represent frames makes it possible to explain phenomena that require structured representations, including analogy, similarity, and metaphor. **Indurkha** further notes that, by virtue of being perceptual, frames provide powerful means of explaining how truly novel features emerge during these phenomena. In a wonderful example, he illustrates how blending a perceptual simulation of sunlight with a perceptual simulation of the ocean's surface the emergent feature of harp strings vibrating back and forth. Indurkha argues compellingly that perceptual symbol systems explain the emergence of such features much more naturally than do amodal symbol systems. In this spirit, he suggests that conceptual blending be called "perceptual blending" to highlight the importance of perceptual representations, and he cites empirical evidence showing that perception enters significantly into metaphor comprehension.

Fauconnier notes that I fail to acknowledge the complexity of simulations in language and thought. He provides a compelling example of a child who uses sugar cubes and matchbooks to simulate cars and buses driving down the street. Clearly, such examples involve complicated simulations that reside on multiple levels, and that map into one another in complex ways. Indeed, this particular example pales in complexity next to Fauconnier and Turner's (1998) examples in which people juxtapose perception of the current situation with past, future, and counterfactual situations. In these examples, people must represent several situations simultaneously for something they say to make sense. To my mind, the fact that the human conceptual system can represent such complicated states-of-affairs illus-

trates how truly remarkable it is. **Brewer's** observation that simulations provide a good account of how scientists represent models further illustrates this ability.

The evolution of the human frontal lobes provides one way to think about people's ability to construct multiple simulations simultaneously and to map between them. Increasingly large frontal lobes may have provided greater inhibitory power to suppress perception and to represent nonpresent situations (cf. Carlson et al. 1998; Donald 1991; 1993; Glenberg et al. 1998). Taking this idea a little further provides leverage in explaining humans' ability to represent multiple nonpresent situations simultaneously. Not only can humans simulate absent situations, we can simulate several absent situations simultaneously and map between them. The possession of a powerful inhibitory and control system in the frontal lobes may well make this possible.

R2.7. Introspection, emotion, and metacognition. Two commentators found my inclusion of introspective processes problematic and unparsimonious (**Newton, Toomela**). Ultimately, they worry about the mechanisms that perceive introspective events. In place of introspection, Newton suggests that we seek subtle proprioceptive events that ground intuitive understandings of introspection. For example, experiences of collateral discharge while executing movements might ground the intuitive concept of *agency*, whereas experiences of eye movements might ground the intuitive concept of *seeing*. Similarly, Toomela suggests that introspective phenomena arise through interactions of more basic systems, such as perception and language.

Although I am sympathetic to these proposals and could imagine them being correct to some extent, I am not at all convinced that it is either necessary or wise to explain all introspection in these ways. Most significantly, people clearly experience all sorts of internal events, including hunger, fatigue, and emotion. Finding ways to ground them in more externally-oriented events seems difficult, and denying their internal experience seems highly counterintuitive.

Another problem for **Newton and Toomela** is explaining the central role that the representation of mental states plays in modern developmental and comparative psychology. Research on theory of mind shows that people represent what they and others are representing (or are not representing) (e.g., Hala & Carpendale 1997). A primary argument to emerge from this literature is that representing mental states is a central ability that distinguishes humans from other species (Tomasello & Call 1997). If we do away with introspection, how do we represent minds in a theory of mind? Also, how do we explain the wide variety of metacognitive phenomena that people exhibit (**Charland**)?

Regarding the problem of what mechanisms perceive representing, I disagree with **Newton and Toomela** that this leads to an infinite regress whereby additional mechanisms to perceive representations must be added to the basic mechanisms that perceive the world. Rather, I suspect that the brain uses one set of mechanisms to perceive both the world and representations. Recall the basic premise of perceptual symbol systems: Top-down simulations of sensory-motor systems represent knowledge, optionally producing conscious states in the process (sect. 2.1). Perhaps the mechanisms that produce conscious experience of sensory-motor systems when driven by bottom-up sensory processing also produce consciousness of the same systems

when driven by top-down activation from associative areas (sect. 2.4.7). No new mechanisms are necessary. The one critical requirement is knowing the source of the information that is driving conscious experience. In psychosis, this awareness is lacking, but most of the time, a variety of cues is typically sufficient, such as whether our eyes are closed (Glenberg et al. 1998), and the vividness of the experience (Johnson & Raye 1981).

I agree strongly with **Charland** that emotion is central to perceptual symbol systems, and that emotion's role in this framework must be developed further. I also agree that there may well be specific brain circuits for processing particular emotions, which later support conceptualizations of these emotions (e.g., circuits for fear and anxiety; Davis 1998). Where I disagree with Charland is in his apparent claim that emotion is a self-contained symbol system. I further disagree with his readings of Damasio (1994) and Lazarus (1991) that perceptual images and appraisal mechanisms belong to an autonomous symbol system for emotion. This is certainly not my reading of their work. On the contrary, I strongly suspect that these cognitive aspects of emotion are provided directly by cognitive – not emotion – mechanisms. Rather than containing its own cognitive system, I suspect that emotion mechanisms complement a separate cognitive system. I hasten to add, however, that these two systems are so tightly coupled that dissociating them creates a distortion of how each one functions alone (Damasio 1994). Also, reenacting states in emotion systems – not amodal systems in a general knowledge store – constitutes the basis of representing emotions symbolically in conceptual processing.

Finally, **Oehlmann** presents two metacognitive phenomena and asks how perceptual symbol systems explain them. First, how do people know that they know the solution to a problem without having to simulate the entire solution? Selective attention and schematization provide one account (sect. 2.2). On simulating a solution, selective attention stores the initial and final states in working memory, dropping the intermediate states. By then switching back and forth between the initial and final states, an abbreviated version of the simulation becomes associated with the entire simulation. On later perceiving the initial conditions for these simulations, both become active, with the abbreviated one finishing much sooner. If the final state in the abbreviated simulation is sufficient to produce a response, waiting for the complete simulation to finish is unnecessary. As this example illustrates, selective attention and schematicity provide powerful editing functions on simulations, just as in productivity, making it possible to rise above complete simulations of experience (sect. 3.1). Again, a perceptual symbol system is not a recording system.

Oehlmann's second problem is how perceptual symbol systems explain performance on the false belief task. Although children can be shown to recognize that another person has a different belief than they do, they nevertheless forget this at times and attribute their own belief to the other person. Oehlmann claims that if beliefs were simply simulated, this pattern of results should not occur. If children can simulate another person's belief correctly on one occasion, why do they not simulate it correctly on all occasions? An obvious explanation is that children at this age are limited in their ability to simulate other people's mental states. Under optimal conditions, they have enough cognitive resources to run different simulations for themselves

and others. Under less than optimal conditions, however, they run out of resources and simply run the same simulation for others as they do for themselves. As children grow older and their ability to simulate other minds automatizes, they almost always have enough resources to simulate differing states of mind, and they rarely fall back on the projective strategy. Indeed, this is similar to Carlson et al.'s (1998) proposal that frontal lobe development is the key resource in performing these tasks. When task demands are too high, children do not have enough resources to inhibit simulations of their own mental state when attempting to simulate the mental states of others.

R2.8. Situated action. I agree with **Glenberg** that if a perceptual symbol system fails to support situated action, it is a limited and misguided enterprise. Whatever form the cognitive system takes, it most certainly evolved to support situated action (e.g., Clark 1997; Glenberg 1997; MacWhinney 1998; Newton 1996). In the target article, I focused primarily on how perceptual symbol systems achieve classic symbolic functions, given that this ability has not been appreciated. In the process, I failed to explore how perceptual symbol systems support situated action. I have attempted to remedy this oversight in a more recent paper, arguing that perceptual simulation is an ideal mechanism for monitoring and guiding situated action (Barsalou, in press). **Freksa et al.** similarly argue that perceptual simulation is well suited for interfacing cognition with action in the world. Because cognition and perception use a common representational system, cognition can be readily linked with perception and action. Internal models have sufficient perceptual fidelity to the physical world to provide accurate and efficient means of guiding interactions with it. Again, though, causal relations, not just resemblance, are important in establishing these relations (sects. 2.4.3, 3.2.8, and R1.2).

R2.9. Development. Two commentaries note that classic developmental theories anticipate the importance of perceptual symbol systems in cognition. **Mitchell & Clement** remind us that Piaget had a lot to say about the development of cognition from sensory-motor processing. **Toomela** reminds us that Vygotsky and Luria had a lot to say, as well. I agree completely and did indeed neglect these theorists in section 4.2, primarily because Nelson's (1996) book, which I did cite, does a good job of describing these previous developmental theories as a springboard for her recent theory.

McCune is concerned about my statement that "the same basic form of conceptual representation remains constant across . . . development, and a radically new form is not necessary." In particular, she notes that children's ability to use representations symbolically develops considerably over the first two years of life. As children mature, frontal lobe development supports increasingly powerful abilities to formulate and manipulate representations offline from perception. I agree completely (sects. R2.6 and R2.7). In the quote above, however, I was not referring to representation in this sense of symbolic activity. Instead, I was arguing that a transition from a perceptual symbol system to an amodal symbol system does not occur over the course of development. Instead, a single perceptual symbol system develops, achieving critical symbolic milestones along the way. I suspect, however, that even the rudiments

of more advanced representational abilities are present very early. As I speculated in section 3.4.3, prenatal infants may compare expectations to sensory experiences as a precursor to the operations that underlie the adult concepts of *truth* and *falsity*. As McCune suggests, however, these early precursors are certainly much more primitive than their later counterparts.

R.3. Abstract concepts

The topic raised most frequently by the commentators was abstract concepts. As I noted in section 3.4, accounting for these concepts in perceptual symbol systems is a controversial and challenging task.

R3.1. Abstract concepts have nothing to do with perception. Arguing that abstract concepts rise above perception, **Landauer** and **Ohlsson** view perception as a nuisance and hindrance to representing these concepts. Neither, however, tells us what abstract concepts are. If they are not about events we perceive, then what are they about? Lacking a compelling account of their content, how can we claim to understand them? By "content" in this discussion, I will mean the cognitive representations of abstract concepts, because this is the type of content at issue here. By no means, however, do I intend to imply that physical referents are unimportant. As described earlier for the causal theory of concepts, I strongly agree that they can be central (sect. R1.2).

In section 3.4.2, I proposed a strategy for representing abstract concepts in perceptual symbol systems: (1) identify the event sequences that frame an abstract concept; (2) specify the relevant information across modalities in these sequences, including introspection; and (3) specify the focal content in these sequences most relevant to defining the concept. Most basically, this strategy requires identifying the entities and events to which an abstract concept refers. That is all. Indeed, it would seem to be the same strategy that one should follow in developing amodal accounts of abstract concepts, as well. Regardless of the specific approach taken, it is essential to identify the content of an abstract concept.

Once one identifies this content, the issue then arises of how to represent it. One possibility is to transduce it into amodal symbols. Another possibility is to simulate it in the spirit of perceptual symbol systems. Regardless, it is essential to identify the content of an abstract concept. Just postulating the concept's existence, or using a language-based predicate to represent it, fails to accomplish anything. There is nothing useful in claiming that TRUE (X) and FALSE (X) represent the abstract concepts of *truth* and *falsity*. Obviously, specifying the content that distinguishes them is essential. Once this has been accomplished, my claim is simply that simulations of this content are sufficient to represent it – there is no need to transduce it into amodal symbols.

Another key issue concerns what types of content are potentially relevant to representing abstract concepts. Throughout the target article, but especially in section 3.4.2, I argued that it is impossible to represent abstract concepts using content acquired solely through sensation of the physical world. On the contrary, it is essential to include content from introspection. Whereas concrete concepts are

typically concerned only with things in the world, abstract concepts are about internal events, as well. Thus, I agree that abstract concepts cannot be explained solely by perception of the external world, and my frequent argument has been that this additional content arises in introspection. Unlike **Landauer** and **Ohlsson**, I have specified where this content is supposed to originate. I further predict that should they attempt to characterize this content themselves, they will ultimately find themselves appealing to introspective events. Most important, I predict that simulations of these events will be sufficient to represent them and that amodal symbols will not be necessary.

R3.2. Abstract concepts of nonexistent entities. Several commentators wondered how perceptual symbol systems could represent something that a person has never experienced. If, as I just suggested, the representation of an abstract concept is the reenactment of its content, how can we ever represent an abstract concept whose referents have never been encountered? How do we represent *the end of time*, *life after death*, *infinite set*, and *time travel* (**Ohlsson**)? How do we represent *electromagnetic field* and *electron* (**Toomela**)?

It is instructive to consider concrete concepts that we have never experienced, such as *purple waterfall*, *cotton cat*, and *battered chair*. Although I have never experienced referents of these concepts, I have no trouble understanding them. Using the productive mechanisms of perceptual symbol systems. I can combine components of past experience to construct novel simulations of these entities (sect. 3.1.2). I can simulate a normal waterfall and transform its color to purple; I can simulate a normal cat and transform its substance to cotton; I can simulate a normal chair and simulate battering it. If I can construct such simulations for concrete concepts, why can I not construct them for abstract concepts?

For *the end of time*, I can begin by simulating the ends of known processes, such as *the end of growth* and *the end of walking*. In all such cases, a process occurring over a stretch of time stops. Applying the same schematic simulation to *time* yields an interpretation of *the end of time*, even though we have not experienced it. Simple-mindedly, we might imagine that all clocks stop running. With a little more sophistication, we might imagine all change in the world ending, with everything freezing in place. Or, we might simply imagine all matter disappearing into nothing. By combining past experience productively, we can simulate events and entities that we have never experienced, regardless of whether they are concrete or abstract (sect. 3.1.2).

If space permitted, I would be happy to provide similar accounts of the other examples raised. However, making even a half-hearted attempt to follow my proposed strategy for representing abstract concepts may well be sufficient for readers to convince themselves that it is quite possible to represent nonexperienced concepts with productive simulations (sect. 3.4.2).

What is perhaps most troubling about raising nonexperienced concepts as a problem for perceptual symbol systems is the lack of awareness that these concepts constitute a challenge for amodal symbol systems, too. How do amodal approaches represent the content of these concepts? Once amodal theorists make the effort to identify this content, I again predict that perceptual symbols will provide a direct and parsimonious account.

R3.3. Metaphor and abstract concepts. In reaction to my argument that metaphor does not form the primary basis of abstract concepts (sect. 3.4.1), **Gibbs & Berg** respond that metaphor does not just elaborate and extend abstract concepts – it is central to their core. In making their case, they cite empirical studies showing that metaphors structure abstract concepts and produce entailments in their representations. My intention was never to underestimate the importance of metaphor in abstract concepts, especially those for which people have no direct experience. I agree that metaphor is important. Instead, my intention was to redress cognitive linguists' failure to consider the importance of direct experience in abstract concepts, and I remain convinced that direct experience is critically important. It remains an open empirical question just how much metaphor structures abstract concepts, and I suspect that the nature of this influence varies widely from concept to concept. Of the greatest importance, perhaps, is establishing a detailed process model of how direct experience and metaphor interact to represent abstract concepts over the course of their acquisition.

R3.4. Situations and abstract concepts. Central to my proposed strategy for representing abstract concepts is identifying the background event sequences that frame them (sect. 3.4). Perhaps the greatest omission in the target article was my failure to cite empirical research on situations that support this conjecture. Much research by Schwanenflugel and her colleagues shows that abstract concepts are harder to process than concrete concepts when background situations are absent, but that both types of concepts are equally easy to process when background situations are present (e.g., Schwanenflugel 1991; Schwanenflugel et al. 1988; Schwanenflugel & Shoben 1983; Wattenmaker & Shoben 1987). Wiemer-Hastings (1998) and Wiemer-Hastings and Graesser (1998) similarly report that sentential contexts are much more predictive of abstract words than of concrete words, further indicating the importance of background situations for abstract concepts. In their commentary here, **Wiemer-Hastings & Graesser** offer excellent examples of how abstract concepts depend on situations for meaning. Understanding why abstract concepts depend so heavily on background situations is likely to be central for developing adequate accounts of them.

R3.5. Representing truth, falsity, and negation. As section 3.4.3 noted, it was never my aim to provide complete accounts of these three concepts. Instead, I presented partial accounts to illustrate a strategy for representing abstract concepts in general (sect. 3.4.2), noting that more complete analyses would be required later.

Landau complains that I did not really specify anything about the content of *truth*. This is indeed quite puzzling, given that the commentators I am about to discuss all believed that I had in fact specified such content. The problem, they thought, was that I had not specified it adequately. Two commentaries suggested that my account of *truth* cannot be distinguished from *match* (**Adams & Campbell**), or from *similar*, *comparable*, and *looks like* (**Mitchell & Clement**). Following the proposed strategy for representing abstract concepts in section 3.4.2, distinguishing these concepts begins by examining the situations in which they occur. In *match*, *similar*, *comparable*, and *looks like*, the relevant situation typically involves compar-

ing two entities or events in the same domain of experience. Thus, one might compare two birds in the world or two imagined plans for a vacation. In none of these cases are two ontologically different things typically compared, such as an imagined perception and an actual one. Furthermore, subtle aspects of comparison distinguish these four concepts from one another. Thus, *match* implies that identical, or nearly identical, features exist between the compared objects, whereas *similar*, *comparable*, and *looks like* imply partially overlapping features. Within the latter three concepts, *similar* allows just about any partial match, *comparable* implies partial matches on aligned dimensions (Gentner & Markman 1997), and *looks like* implies partial matches in vision.

In contrast, the sense of *truth* in section 3.4.3 and Figure 7a specifies that an agent attempts to map a mental simulation about a perceived situation into the perceived situation, with the simulation being true if it provides an adequate account of the perceived situation. Clearly, the components of this account differ from those just described for *match*, *similar*, *comparable*, and *looks like*. In *truth*, one entity (a mental simulation) is purported to be about a second entity (a perceived situation). In none of the other four concepts do the two entities being compared reside in different modalities, nor is one purported to be about the other. As we shall see, more content must be added to this account of *truth* to make it adequate. Nevertheless, enough already exists to distinguish it from these other concepts.

Once one establishes such content, the key issue, of course, is deciding how to represent it. Following the amodal view, we could attempt to formulate complicated predicate calculus or programming expressions to describe it. Alternatively, we could attempt to identify the direct experience of this content in the relevant situations and argue that simulations of these experiences represent the content. Thus, over many occasions of assessing *truth*, a simulator develops that can construct specific simulations of what it is like to carry out this procedure.

Siebel notes an omission in my account of *truth* in Figure 7a. How does it explain errors in mapping simulations to perceived situations? If I simulate a donkey and mistakenly bind it to a perceived horse, do I believe it true that the perceived entity is a donkey? If I have no disconfirming evidence, yes! Of course, this does not mean that my simulation is actually true, it simply means that I have concluded incorrectly that it is. Indeed, people often believe that an expectation is true, even before they compare it to the perceived world (Gilbert 1991).

Where the account of *truth* in Figure 7a must be extended is to account for the discovery of disconfirming evidence. Following the proposed strategy for representing abstract concepts (sect. 3.4.2), simulations of disconfirmation must be added. Because many kinds of disconfirmation are possible, many event sequences are necessary. For example, further perception of the horse might activate additional visual features, which in turn activate the simulator for *horse*. After comparing simulations for *horse* and *donkey* to the perceived entity, I might decide that the original donkey simulation does not fit as well as the subsequent horse simulation, therefore changing what I believe to be true. Alternatively, another agent with more expertise than me might claim that the perceived entity is a horse, so that I attribute falsity to my original simulation and attribute truth to a new simulation from my horse simulator. By in-

corporating schematic simulations of disconfirmation into the simulator for *truth*, the account in Figure 7a can be extended to handle disconfirmation. It is simply a matter of implementing the strategy in section 3.4.2 to discover the relevant content.

Ohlsson notes another omission in my account of *truth*. Rather than applying my proposed strategy to see if it yields the necessary content, however, he simply concludes that it cannot. The omission he notes is that lack of a fit does not necessarily mean a simulation is false. In his example, I see a cat on a mat in my office. While I am turned away, someone pulls the mat outside my office with the cat on it. When I turn around, the cat and the mat are gone. Ohlsson claims that my account of *falsity* in Figure 7b implies that it is false that the cat is on the mat. Actually, though, this is not what my account would say. As I stated repeatedly in the target article, a simulation is purported to be about a perceived situation (sect. 3.4.3). Thus, in Ohlsson's example, my simulation of a cat on a mat is purported to be about my office. When this simulation no longer matches the perceived situation, one reasonable conclusion is that it is false that there is a cat on a mat *in my office*, not Ohlsson's conclusion that it is false that the cat is on the mat. If the critical question were whether the cat is on the mat *anywhere*, then I would have to compare a simulation of a cat on a mat to many perceived situations until I found one containing the mat.

This adds another important layer of content to my accounts of *truth* and *falsity* in Figure 7: On some occasions in which these concepts are used, it may be necessary to compare a simulation to more than one perceived situation. Only after failing to find a satisfactory fit in any relevant situation does a conclusion about falsity follow, assuming that a match in a single situation suffices for *truth* (i.e., in a universally quantified claim, a match would be necessary in every situation). The same solution handles **Ohlsson's** other example about aliens. We do not conclude that there are no aliens in the universe after examining only one situation. Instead, we have to examine many possible situations before reaching this conclusion. Again, however, applying the strategy in section 3.4.2 yields the necessary content. By examining the appropriate situations, the relevant content can be discovered and added to the respective simulators.

Fauconnier notes that a near miss is a better example of *falsity* than a far miss. On perceiving a balloon above a cloud, for example, a simulation of a balloon under a cloud is a better example of a false simulation than is a simulation of a typewriter on a desk. Such a prediction is borne out by an increasingly large literature which shows that high similarity counterintuitively produces high dissimilarity (Gentner & Markman 1997). The more alignable two representations are, the more similar they are, *and* the more different they are. Thus, I suspect that Fauconnier's prediction has more to do with how people compute fits between representations than with my accounts of *truth* and *falsity* per se. His prediction not only applies to these concepts but to all of the related concepts discussed earlier, such as *match*, *similar*, *comparable*, and *looks alike*.

Finally, **Wiemer-Hastings & Graesser** explore my point in section 3.4.3 that *truth* is a polysemous concept. In their commentary, they explore a variety of senses that I did not attempt to cover in Figure 7a. Besides highlighting the fact that most abstract concepts are highly polysemous, they raise the further issue that any theory of knowledge must

differentiate the senses of an abstract concept. Using the concept of *idea*, Wiemer-Hastings & Graesser illustrate how situations may well provide leverage on this task. I suspect that the same strategy is also likely to provide leverage on the different senses of *truth* that they raise, such as the truth of a single utterance, the truth of all a person's utterances, and scientists trying to discover the truth. As just described for *falsity*, differences in how simulations are assessed against perceived situations may distinguish the first two senses. In the first sense, a single simulation for one particular utterance must be compared to a single perceived situation. In the second sense, a simulation for each claim a person makes must be compared to each respective situation. To account for the scientific sense of *truth*, a completely new set of situations from the scientific enterprise is required, including the formulation of hypotheses, the conducting of research, the assessment of the hypotheses against data, and so forth. Only within this particular set of experiences does the scientific sense of *truth* make sense. Again, specifying the meaning of an abstract concept requires searching for the critical content in the relevant background situations.

R3.6. How do you represent X? Various commentators wondered how perceptual symbols could represent all sorts of other abstract concepts. These commentaries fall into two groups. First, there was the commentator who followed my strategy in section 3.4.2 for representing abstract concepts and discovered that it does indeed provide leverage on representing them. Second, there were the commentators who complained that perceptual symbol systems fail to represent particular abstract concepts, yet who apparently did not try this strategy. At least these commentators do not report that this strategy failed on being tried.

First, consider the commentator who tried the strategy. **Hurford** notes that the target article failed to provide an account of the concept *individual*. On examining relevant situations that contain familiar individuals, Hurford induced an account of this concept: A unique individual exists whenever we fail to perceive anything exactly like it simultaneously in a given situation. Thus, we conclude that our mother is a unique individual because we never perceive anyone exactly like her simultaneously. For the same reason, we conclude that the sun and moon are unique individuals. As Hurford further notes, mistakes about the multiple identities of the same individual follow naturally from this account. Because we perceive Clark Kent and Superman as each being unique, we never realize that they are the same individual (the same is true of the Morning Star and the Evening Star). Only after viewing the transition of one into the other do we come to believe that they are the same individual. Like my account of *truth*, Hurford's account of *individual* may need further development. Regardless, it illustrates the potential of the strategy in section 3.4.2 for discovering and representing the content of abstract concepts.

Finally, consider the commentators who doubted that perceptual symbol systems could represent particular abstract concepts, apparently without trying the strategy in section 3.4.2. **Adams & Campbell** expressed skepticism about *chiliagons* and *myriagons*; **Brewer** expressed skepticism about *entropy*, *democracy*, *evolution*, and *because*; **Mitchell & Clement** expressed skepticism about *can*, *might*, *electricity*, *ignorant*, and *thing*; **Ohlsson** expressed

skepticism about *health care system*, *middle class*, *recent election*, and *internet*. Again, if there were no shortage of space, I would be happy to illustrate how the proposed strategy in section 3.4.2 might represent these concepts. Again, the content of an abstract concept must be specified, regardless of whether it is to be represented perceptually or amodally, and the proposed strategy is a useful means of identifying it. Furthermore, I remain convinced that simulations of this content are sufficient to represent it, and that amodal redescriptions are unnecessary.

R4. Language

Another frequently raised topic was the relation between language and perceptual symbol systems.

R4.1. Linguistic forms per se do not carry meaning. I continue to be amazed by how often sophisticated researchers believe that language per se carries meaning. Most commonly, this belief has to do with internal speech. As someone listens to imagined speech in introspection, the speech is presumed to carry meaning as the person "thinks in language." As **Fauconnier** notes, however, language does not carry meaning directly. Instead, language is a powerful means of controlling the online construction of conceptualizations (sect. 2.7). One can hear a foreign language thousands of times, but unless the words become linked to their referents in experience, they remain meaningless. The words do not inherently carry meaning.

In arguing that we think in language when processing abstract concepts, Jorion strikes me as being under the illusion that language carries meaning (a similar criticism of **Landauer** was presented earlier in sect. R1.4). In a series of examples, **Jorion** argues that concepts are represented in words, not in percepts. The fact that some people are highly conscious of internal speech, though, does not entail that words are the causal forces in their thinking. As most cognitive scientists agree, the bulk of the action in cognition is unconscious, and a lot more is happening than conscious inner speech. If language is so central to thought, how can aphasics think, and how can **Lowenthal's** subjects learn? How can people be thinking of something and not know the word for it? Where do all the inferences in comprehension come from that are not expressed directly in linguistic forms?

R4.2. Language in knowledge acquisition. I hasten to add that language most certainly plays critical roles in the acquisition and use of knowledge. For example, much research on conceptual development has shown that words provide important signals for developing new concepts (e.g., Gelman et al. 1998; Gelman & Tardif 1988; Markman 1989). On hearing a new word, children often infer that it refers to a new set of entities that share important commonalities. **Hurford** similarly notes how linguistic elements may signal the presence of individuals in domains of discourse, perhaps assisting in the acquisition of *individual*. **Toomela** similarly suggests that language is central to forming abstract concepts. Finally, **Langacker's** examples of construal illustrate how the word associated with a perceived entity determines its conceptualization, with different words producing different conceptualizations (sects. 2.4.7 and 3.2.6).

Lowenthal provides an example to the contrary, show-

ing that people can acquire new knowledge without language. Even when lesions in language areas of the brain disrupt linguistic processing, people still learn new concepts and skills. As I have argued throughout this reply, such demonstrations strongly question the grounding of knowledge in language. Again, it by no means follows that language is normally irrelevant.

R4.3. Language in human evolution. In section 4.2., I suggested that human intelligence may have diverged from nonhuman primates with the evolution of language, citing Donald (1991; 1993). **Gabora**, however, cites a passage from Donald in which he states that some sort of conceptual foundation, which may have been mimetic culture, must have preceded speech. In other words, the divergence must have occurred after the evolution of some more basic representational ability, not after the evolution of language. I stand corrected. Gabora has indeed described Donald's correct position. Ultimately, though, my point was that humans evolved the means to control joint simulations about nonpresent situations. By developing the ability to represent nonpresent situations jointly in the past and future, humans increased their fitness substantially. Using the body to act out nonpresent situations (mimesis) may well have been the first stage of this evolution, with language building on it later. Again, my point was that developing the ability to simulate nonpresent situations constituted a major development in human evolution.

R4.4. Cognitive linguistics. As noted by **Fauconnier** and **Langacker**, cognitive linguistics and perceptual symbol systems share many common assumptions. Most importantly, they assume that conceptualizations are grounded in the experiential systems of the brain and body. What I tried to show in the target article is that both approaches belong to a much larger tradition that stretches back over two millennia (sects. 1.1 and 1.3). Embodied approaches to cognition are hardly novel and only appear so in the context of the twentieth century. What is new and exciting is the reinvention of this approach in the contexts of modern cognitive science and neuroscience. Contrary to **Landau** and **Ohlsson's** claim that perceptual symbol systems offer nothing new, it and other current formulations depart from earlier formulations through the incorporation of modern findings, insights, and tools.

Quite a few people have suggested to me informally that cognitive linguistics and perceptual symbol systems are complementary, lying at different levels of analysis. Although cognitive linguists develop accounts of semantics and grammar in terms of experiential constructs, these accounts often fail to map clearly into well-established cognitive and neural mechanisms. One goal of perceptual symbol systems was to provide an account of how a conceptual system could be grounded in these mechanisms. Showing that concepts can be grounded in perception and movement not only supports the assumptions of cognitive linguists, it provides a layer of theoretical machinery that may be useful to their accounts. Conversely, cognitive linguists' analyses of myriad linguistic structures strongly suggest various conceptual structures that we are well advised to look for in cognition and neuroscience. As usual, the interplay between levels is not only stimulating but probably necessary for success in achieving the ultimate goals of these fields.

R4.5. Perceptual simulation in comprehension. Because amodal theories of knowledge have dominated accounts of language comprehension, evidence that perceptual simulation underlies comprehension instead is significant. If the meanings of texts are grounded in simulations, one of the strongest arguments for the amodal view is in danger. **Zwaan et al.** review a variety of findings which suggest that simulations do indeed underlie comprehension. Zwaan et al. also suggest a variety of future projects that could explore this issue further. In section 4.1.6 of the target article, a variety of additional studies are cited that support embodied comprehension, and Barsalou (in press) provides an evolutionary account of why comprehension might have become grounded in perceptual simulation. Clearly, much additional research is needed to address this hypothesis thoroughly, but existing evidence suggests that it may well be true.

R5. Computational accounts

As I mentioned throughout the target article, developing computational accounts of perceptual symbol systems is important (e.g., sects. 2 and 5). A number of commentators agreed, and had a variety of things to say about this issue.

R5.1. Functional specifications versus mechanisms. Several commentators suggested that the target article provided only functional specifications of perceptual symbol systems and failed to say anything at all about the underlying mechanisms (**Adams & Campbell; Dennett & Viger; Schwartz et al.**). As I acknowledged in the target article, I did not provide a full-blown computational account of perceptual symbol systems. Nevertheless, I certainly did have something to say about mechanisms. Throughout the target article, I discussed the viability of Damasio's (1989) convergence zones, arguing that they provide an ideal mechanism for implementing perceptual symbols (sects. 2.1, 2.2, 2.3, 2.4.6, 2.4.7, and 4.3). In a convergence zone, one layer of units represents information in a sensory-motor domain, and a second layer in an association area captures patterns of activation in the first layer. By activating a pattern of units in the associative area, a sensory-motor representation is reenacted. The mechanics of such systems are basic lore in connectionism, and in citing convergence zones frequently throughout the target article, I was pointing to an obvious and well-known mechanism for implementing perceptual symbols.

During discussions of symbol formation (sects. 2.2 and 2.3), I suggested a second line of mechanisms that are central to perceptual symbol systems. Specifically, I suggested that the well-known mechanisms of attention and memory extract and store components of experience for later use as symbols. Although I did not provide detailed accounts of how these mechanisms work, I did indeed incorporate them centrally. Furthermore, it would be possible to draw on detailed accounts in the literature to develop my proposed account of how perceptual symbols are extracted from experience.

The discussions of frames similarly proposed mechanisms that underlie the formation of simulators. In section 2.5 and Figure 3, I sketched an account of the associative mechanisms that underlie frames. In section 2.5.1, I defined the components of frames and described how they could be implemented in perceptual symbol systems.

In my opinion, I developed the mechanisms of convergence zones, symbol extraction, and frames enough to show that they can potentially achieve the criteria of a fully functional conceptual system. By combining attention, memory, and convergence zones, I illustrated how elemental perceptual symbols could become established in memory. By showing how these elemental symbols combine to form frames and simulators, I illustrated how concepts and types could be represented. By showing how simulators combine with each other and with perceived situations, I illustrated how they could explain productivity and propositions. The basic strategy of my argument was to show how preliminary sketches of convergence zones, symbol extraction, and frames hold promise for constructing a fully functional conceptual system.

I was the first to note that these mechanisms must be specified further before this argument is fully compelling (sect. 5). Nevertheless, I defend my decision to work at a relatively high level initially. Building a detailed implementation that ultimately constitutes a significant contribution requires this sort of preliminary analysis. If we can first agree on what an implementation should do, and on what basic mechanisms it should have, we will save a lot of time and energy. For this reason, I urge patience before firing up the compiler, as well as appreciation for discussions of what implementations should accomplish.

R5.2. Connectionism. Several commentators came away with the mistaken impression that I am not a fan of connectionism (**Gabora; Schwartz et al.**). These commentators appeared to take my criticism of amodal neural nets in section 1.2 as a criticism of connectionism in general. Again, however, my criticism was of those particular systems in which a hidden layer of conceptual units acts as an amodal symbol system, being separate from and arbitrarily connected to a layer of perceptual units. As I stated in that section, my criticism does not apply to all connectionist theories. Furthermore, in section 2.2.1, I described explicitly how perceptual symbols are dynamical attractors in neural nets. Throughout section 2.5, I similarly described how the mechanisms of frames operate according to associative and dynamical principles. Thus, I find it difficult to understand how I could be construed as not recognizing the virtues of connectionism for implementing perceptual symbol systems. I do! Again, I strongly endorse connectionist schemes in which a common population of units represents information in perception and knowledge, and I view Damasio's (1989) convergence zones as an appropriate architecture for implementing this. Trehub (1977; 1991) suggests another architecture that may be useful for implementing perceptual symbol systems, as well.

R5.3. Implementing perceptual symbol systems. In this final section, I suggest four aspects of perceptual symbol systems that strike me as particularly important and challenging to implement. First, and perhaps foremost, implementing mechanisms that represent space and time is absolutely essential. Intuitively, these mechanisms implement a virtual stage or platform on which simulations run. It is the "virtual space" or Kantian manifold in which simulated events unfold over time. Perhaps it is essentially the same system that represents the perception of actual events in the world as they unfold. Regardless, perceptual symbol systems cannot be implemented adequately until it is pos-

sible to run simulations either in isolation or simultaneously with perception. Note that such a system need not be topographically organized with respect to space and time, although it could be. Functional representations of space and time may well be sufficient.

As several commentators suggested, representing space is essential to implementing structured representations and productivity (**Edelman & Breen; Markman & Dietrich**). As I noted earlier (sect. R2.5), representing time is no less important. Figuring out how to implement space and time will not only be critical for running simulations, it will also be critical for implementing classic symbolic functions. To create structured representations, it will be necessary to define space-time regions that can be specialized.

Second, it is essential to develop detailed mechanisms for implementing frames (**Adams & Campbell; Dennett & Viger; Markman & Dietrich**). Again, the formulation in section 2.5 and Figure 3 sketches an initial approach. As I noted in Note 14, however, no account of space was provided, much less an account of time. Until mechanisms for representing space and time are implemented, it will be difficult, if not impossible, to implement a viable account of frames. In particular, it will not be possible to implement the basic frame properties of conceptual relations, bindings, and recursion, which rely on the specialization of space-time regions (sects. 2.5.1, 3.1, 3.2, 3.4.4, and R2.5). Once a system for representing space and time is implemented, it can be used to integrate the perceptual symbols extracted from different instances of a category into a common frame, as Figure 3 illustrated.

Third, once a system for representing frames is in place, it will be critical to specify how a particular simulation is assembled from the information in a frame. Again, mechanisms for representing space and time are likely to be important. Because a space-time region typically contains only a single entity or event, this constraint blocks most information in a frame from becoming active. Figuring out how to implement this constraint will be important, as will figuring out how to control the sketchiness of simulations. In highly schematic simulations, many space-time regions may be blank or filled minimally, whereas in highly detailed simulations, most space-time regions may be filled richly. Finally, it will be necessary to figure out how to maintain coherence within a simulation, so that incompatible content does not typically become active simultaneously across different space-time regions. The standard mechanisms of constraint satisfaction are likely to be useful.

Finally, the control of attention will be absolutely critical (**Hochberg, Toomela**). To extract components of experience during the symbol formation process, it will be necessary to specify the focus of attention (sect. 2.2). To implement productivity during the construction of hierarchical simulations, specifying the control of attention will again be necessary so that space-time regions can be specialized selectively (sects. 3.1 and 3.4.4). To implement the formation of propositions, attention will have to focus on regions of perception so that the relevant simulators can become bound to them (sect. 3.2). As described in sections 1.4 and 4.1.3, the ability to attend selectively to space-time regions frees a theory of representation from the constraints of a recording system and allows it to become a conceptual system.

I am highly sympathetic to **Hochberg's** and **Toomela's** proposal that situated action is likely to be central in guid-

ing attention. However, mechanisms for representing space and time are perhaps more basic. In the absence of spatio-temporal reference systems, it is impossible to compute an attentional trajectory. Guiding attention to achieve a goal presupposes the ability to guide attention through space and time. Thus, for still another reason, developing space-time mechanisms is essential to implementing perceptual symbol systems. Once these mechanisms are in place, it will be possible to formulate mechanisms that guide attention selectively over space-time regions to process them.

Perhaps I am naive about computational modeling, but the four aspects of perceptual symbol systems whose implementation I just outlined strike me as particularly challenging. To my knowledge, nothing “off the shelf” is available for fully implementing these functions, nor are they likely to be implemented in a few weeks’ work. So please forgive me for not yet having implemented the theory in the target article! Perhaps some readers will agree that sketching the theory’s basic mechanisms and illustrating their potential for achieving a fully functional conceptual system constitutes a useful contribution.

R6. Conclusion

The way I see it, three basic approaches to knowledge exist in modern cognitive science and neuroscience: (1) classic representational approaches based on amodal symbols, (2) statistical and dynamical approaches such as connectionism and neural nets, and (3) embodied approaches such as classic empiricism, cognitive linguistics, and situated action. If I were to say that there is value in all three of these approaches, I might be viewed as simply being diplomatic. On the contrary, I believe that each of these approaches has discovered something fundamentally important about human knowledge, which the other two have not. Classic amodal approaches have discovered the importance of structured representations, productivity, and propositions. Statistical approaches have discovered the importance of generalization, partial matches, adaptation, frequency effects, and pattern completion. Embodied approaches have discovered the importance of grounding knowledge in sensory-motor mechanisms and the significance of environments and action for cognition. What I have tried to do in formulating perceptual symbol systems is to integrate the positive contributions of all three approaches. Regardless of whether my particular formulation succeeds, I predict that whatever approach ultimately does succeed will similarly attempt to integrate representation, statistical processing, and embodiment.

ACKNOWLEDGMENT

This work was supported by National Science Foundation Grants SBR-9421326 and SBR-9796200.

References

Letters “a” and “r” appearing before authors’ initials refer to target article and response, respectively.

- Adams, F. & Aizawa, K. (1994) Fodorian semantics. In: *Mental representation*, ed. S. Stich & T. Warfield. Blackwell. [FA]
- Ahn, W. & Bailenson, J. (1996) Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *Cognitive Psychology* 31:82–123. [aLWB]
- Ahn, W., Kalish, C. W., Medin, D. L. & Gelman, S. A. (1995) The role of covariation versus mechanism information in causal attribution. *Cognition* 54:299–352. [aLWB]
- Aizawa, K. (1997) Explaining systematicity. *Mind and Language* 12:115–36. [FA]
- Alpern, M. (1972) Eye movements. In: *Visual psychophysics*, ed. D. Jameson & L. M. Hurvich. Springer-Verlag. [AT]
- Anderson, J. R. (1978) Arguments concerning representations for mental imagery. *Psychological Review* 85:249–77. [aLWB]
- (1993) *Rules of the mind*. Erlbaum. [aLWB]
- Anderson, J. R. & Bower, G. H. (1973) *Human associative memory*. Winston. [aLWB]
- Aristotle (1961) *De anima, Books II and III*, trans. D. W. Hamlyn. Oxford University Press. (Original work published in 4th century B.C.). [aLWB]
- Baddeley, A. D. (1986) *Working memory*. Clarendon. [aLWB]
- Bailey, D., Feldman, J., Narayanan, S. & Lakoff, G. (1997) Embodied lexical development. *Proceedings of the Nineteenth Annual Meeting of the Cognitive Science Society*. Erlbaum. [aLWB]
- Baillargeon, R. (1995) A model of physical reasoning in infancy. In: *Advances in infancy research, vol. 9*. Ablex. [aLWB]
- Barkowsky, T. & Freksa, C. (1997) Cognitive requirements on making and interpreting maps. In: *Spatial information theory: A theoretical basis for GIS*, ed. S. Hirtle & A. Frank. Springer. [CF]
- Barsalou, L. W. (1983) Ad hoc categories. *Memory and Cognition* 11:211–27. [aLWB]
- (1985) Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11:629–54. [aLWB]
- (1987) The instability of graded structure in concepts. In: *Concepts and conceptual development: Ecological and intellectual factors in categorization*, ed. U. Neisser. Cambridge University Press. [aLWB]
- (1989) Intra-concept similarity and its implications for inter-concept similarity. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony. Cambridge University Press. [aLWB]
- (1991) Deriving categories to achieve goals. In: *The psychology of learning and motivation: Advances in research and theory, vol. 27*. Academic Press. [aLWB]
- (1992) Frames, concepts, and conceptual fields. In: *Frames, fields, and contrasts: New essays in lexical and semantic organization*, ed. A. Lehrer & E. F. Kittay. Erlbaum. [arLWB, ABM]
- (1993) Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In: *Theories of memories*, ed. A. C. Collins, S. E. Gathercole & M. A. Conway. Erlbaum. [arLWB, WFB]
- (1995) Storage side effects: Studying processing to understand learning. In: *Goal-driven learning*, ed. A. Ram & D. Leake. MIT Press. [aLWB]
- (in press) Language comprehension: Archival memory or preparation for situated action? *Discourse Processes*. [rLWB]
- Barsalou, L. W. & Hale, C. R. (1993) Components of conceptual representation: From feature lists to recursive frames. In: *Categories and concepts: Theoretical views and inductive data analysis*, ed. I. Van Mechelen, J. Hampton, R. Michalski & P. Theuns. Academic Press. [arLWB]
- Barsalou, L. W., Huttenlocher, J. & Lamberts, K. (1998) Basing categorization on individuals and events. *Cognitive Psychology* 36:203–72. [aLWB]
- Barsalou, L. W. & Prinz, J. J. (1997) Mundane creativity in perceptual symbol systems. In: *Conceptual structures and processes: Emergence, discovery, and change*, ed. T. B. Ward, S. M. Smith & J. Vaid. American Psychological Association. [aLWB]
- Barsalou, L. W., Solomon, K. O. & Wu, L. L. (in press) Perceptual simulation in conceptual tasks. In: *Cultural, typological, and psychological perspectives in cognitive linguistics: The proceedings of the 4th Conference of the International Cognitive Linguistics Association, vol. 3*, ed. M. K. Hiraga, C. Sinha & S. Wilcox. John Benjamins. [arLWB]
- Barsalou, L. W., Yeh, W., Luka, B. J., Olseth, K. L., Mix, K. S. & Wu, L.-L. (1993) Concepts and meaning. In: *Chicago Linguistics Society 29: Papers from the parasessions on conceptual representations, vol. 2*, ed. K. Beals, G. Cooke, D. Kathman, K. E. McCullough, S. Kita & D. Testen. Chicago Linguistics Society. [aLWB]
- Barwise, J. & Etchemendy, J. (1990) Information, infons, and inference. In: *Situation theory and its application*, ed. R. Cooper, K. Mukai & J. Perry. University of Chicago Press. [aLWB]
- (1991) Visual information and valid reasoning. In: *Visualization in mathematics*, ed. W. Zimmerman & S. Cunningham. Mathematical Association of America. [aLWB]
- Barwise, J. & Perry, J. (1983) *Situations and attitudes*. MIT Press. [aLWB]
- Bassok, M. (1997) Object-based reasoning. In: *The psychology of learning and motivation, vol. 37*, ed. D. L. Medin. Academic Press. [aLWB]
- Bates, E., Thal, D., Trauner, D., Fenson, J., Aram, D., Eisele, J. & Nass, R. (1997)

- From first words to grammar in children with focal brain injury. *Developmental Neuropsychology* 13(3): 275–343. [LM]
- Bear, M. F., Connors, B. W. & Paradiso, M. A. (1996) *Neuroscience: Exploring the brain*. Williams & Wilkins. [aLWB]
- Berendt, B., Barkowsky, T., Freksa, C. & Kelter, S. (1998) Spatial representation with aspect maps. In: *Spatial cognition - an interdisciplinary approach to representing and processing spatial knowledge*, ed. C. Freksa, C. Habel & K. F. Wender. Springer. [CF]
- Berkeley, G. (1710/1982) *A treatise concerning the principles of human knowledge*. Hackett. (Original work published in 1710). [aLWB]
- Biederman, I. (1987) Recognition-by-components: A theory of human image understanding. *Psychological Review* 94:115–47. [aLWB, ABM]
- Biederman, I. & Gerhardstein, P. C. (1993) Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance* 19:1162–82. [aLWB]
- Black, J. B., Turner, T. J. & Bower, G. H. (1979) Point of view in narrative comprehension, memory, and production. *Journal of Verbal Learning and Verbal Behavior* 18:187–98. [aLWB]
- Black, M. (1962a) Metaphor. In: *Models and metaphors*, ed. M. Black. Cornell University Press. [BI]
- (1962b) Models and archetypes. In: *Models and metaphors*, ed. M. Black. Cornell University Press. [WFB]
- (1979) More about metaphor. In: *Metaphor and thought*, ed. A. Ortony. Cambridge University Press. [BI]
- Block, N. (1983) Mental pictures and cognitive science. *Philosophical Review* 93:499–542. [aLWB]
- Bloom, P. & Markson, L. (1998) Intention and analogy in children's naming of pictorial representations. *Psychological Science* 9:200–204. [aLWB]
- Borsook, D., Becerra, L., Fishman, S., Edwards, A., Jennings, C. L., Stojanovic, M., Papinicolas, L., Ramachandran, V. S., Gonzalez, R. G. & Breiter, H. (1998) Acute plasticity in the human somatosensory cortex following amputation. *Neuroreport* 9:1013–17. [SE]
- Bottini, G., Corcoran, R., Sterzi, R., Paulesu, E., Schenone, P., Scarpa, P., Frackowiak, R. S. J. & Frith, C. D. (1994) The role of the right hemisphere in the interpretation of figurative aspects of language: A positron emission tomography activation study. *Brain* 117:1241–53. [BI]
- Braitenberg, V. (1978) Cortical architectonics: General and areal. In: *Architectonics of the cerebral cortex*, ed. M. A. Brazier & H. Petsche. Raven Press. [DAS]
- Brandimonte, M. A., Schooler, J. W. & Gabbino, P. (1997) Attenuating verbal overshadowing through color retrieval cues. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:915–31. [aLWB]
- Bransford, J. D. & Johnson, M. K. (1973) Considerations of some problems of comprehension. In: *Visual information processing*, ed. W. G. Chase. Academic Press. [aLWB]
- Brewer, W. F. (1975) Memory for ideas: Synonym substitution. *Memory and Cognition* 3:458–64. [WFB]
- (in press) Scientific theories and naive theories as forms of mental representation: Psychologism revived. *Science and Education*. [WFB]
- Brewer, W. F. & Pani, J. R. (1983) The structure of human memory. In: *The psychology of learning and motivation: Advances in research and theory, vol. 17*, ed. G. H. Bower. Academic Press. [rLWB]
- (1996) Reports of mental imagery in retrieval from long-term memory. *Consciousness and Cognition* 5:265–87. [WFB]
- Brewer, W. F. & Treyens, J. C. (1981) Role of schemata in memory for places. *Cognitive Psychology* 13:207–30. [WFB]
- Brooks, L. R. (1978) Nonanalytic concept formation and memory for instances. In: *Cognition and categorization*, ed. E. Rosch & B. B. Lloyd. Erlbaum. [aLWB]
- Bruner, J. S. (1966) On cognitive growth. In: *Studies in cognitive growth*, ed. J. S. Bruner, R. R. Olver & P. M. Greenfield. Wiley. [FL]
- Buckner, R. L., Petersen, S. E., Ojemann, J. G., Miezin, F. M., Squire, L. R. & Raichle, M. E. (1995) Functional anatomical studies of explicit and implicit memory retrieval tasks. *Journal of Neuroscience* 15:12–29. [aLWB]
- Burbaud, P. (1988) Traitement numérique et latéralisation corticale (IRMF). First meeting of the Société d'anatomie fonctionnelle cérébrale, Paris, 1988 (personal communication). [FL]
- Burbaud, P., Degreze, P., Lafon, P., Franconi, J.-M., Bouligand, B., Bioulac, B., Caille, J.-M. & Allard, M. (1995) Lateralization of prefrontal activation during internal mental calculation: A functional magnetic resonance imaging study. *APSTRACTS 2:0263N. Journal of Neurophysiology* 74(5):2194–2200. [FL]
- Burgess, C. & Chiarello, C. (1996) Neurocognitive mechanisms underlying metaphor comprehension and other figurative language. *Metaphor and Symbolic Activity* 11(1):67–84. [BI]
- Burgess, C. & Lund, K. (1997) Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes* 12:177–210. [aLWB]
- Cacciari, C. & Levorato, M. C. (1994) “See,” “look,” and “glimpse”: What makes the difference? Paper presented at the Meeting of the Psychonomic Society, St. Louis, MO. [aLWB]
- Cann, R. (1993) *Formal semantics*. Cambridge University Press. [JRH]
- Carey, S. (1985) *Conceptual change in childhood*. MIT Press. [aLWB]
- Carlson, S. M., Moses, L. J. & Hix, H. R. (1998) The role of inhibitory processes in young children's difficulties with deception and false belief. *Child Development* 69:672–91. [rLWB]
- Carnap, R. (1958) *Introduction to symbolic logic and its applications*. Dover. [JRH]
- Carroll, L. (1960) *The annotated Alice: Alice's adventures in wonderland and Through the looking glass*. Bramhall House. [aLWB]
- Chalmers, D. J., French, R. M. & Hofstadter, D. R. (1992) High-level perception, representation, and analogy: A critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence* 4:189–211. [BI]
- Chambers, D. & Reisberg, D. (1992) What an image depicts depends on what an image means. *Cognitive Psychology* 24:145–74. [aLWB]
- Charland, L. C. (1995) Emotion as a natural kind: Towards a computational foundation for emotion theory. *Philosophical Psychology* 8(1):59–84. [LCC]
- (1996) Feeling and representing: Computational theory and the modularity of affect. *Synthese* 105:273–301. [LCC]
- (1997) Reconciling cognitive and perceptual theories of emotion. *Philosophy of Science* 64:555–79. [LCC]
- Cheng, P. W. & Novick, L. R. (1992) Covariation in natural causal induction. *Psychological Review* 99:365–82. [aLWB]
- Chomsky, N. (1957) *Syntactic structures*. Mouton. [arLWB]
- (1975) *Reflections on language*. Pantheon Books. [FL]
- Churchland, P. S. (1986) *Neurophilosophy: Toward a unified science of the mind/brain*. MIT Press. [aLWB]
- Clark, A. (1997) *Being there: Putting brain, body, and world together again*. MIT Press. [arLWB]
- Cohen, L. B. (1991) *Infant attention: An information processing approach*. In: *Newborn attention: Biological constraints and the influence of experience*, ed. M. J. Salomon Weiss & P. R. Zelazo. Ablex. [aLWB]
- Cohen, P. R., Atkin, M. S., Oates, T. & Beal, C. R. (1997) Neo: Learning conceptual knowledge by interacting with an environment. *Proceedings of the First International Conference on Autonomous Agents*. ACM Press. [aLWB]
- Compton, B. J. (1995) *Retrieval strategies in memory-based automaticity*. Doctoral dissertation, University of Illinois, Champaign-Urbana. [aLWB]
- Conway, M. A. (1990) *Autobiographical memory: An introduction*. Open University Press. [aLWB]
- Cornsweet, T. N. (1970) *Visual perception*. Academic Press. [AT]
- Coulson, S. (1997) *Semantic leaps: Frame-shifting and conceptual blending*. Doctoral dissertation, University of California, San Diego. [GF]
- Cowey, A. & Stoerig, P. (1991) The neurobiology of blindsight. *Trends in Neuroscience* 14:140–45. [aLWB]
- Craik, F. I. M. & Lockhart, R. S. (1972) Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior* 11:671–84. [aLWB]
- Craik, K. (1943) *The nature of explanation*. Cambridge University Press. [aLWB]
- Crammond, D. J. (1997) Motor imagery: Never in your wildest dreams. *Trends in Neuroscience* 20:54–57. [aLWB]
- Craver-Lemley, C. & Reeves, A. (1997) Effects of imagery on vernier acuity under conditions of induced depth. *Journal of Experimental Psychology: Human Perception and Performance* 23:3–13. [aLWB]
- Cummins, R. (1996) *Representations, targets, and attitudes*. MIT Press. [aLWB]
- Damasio, A. R. (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33:25–62. [arLWB]
- (1994) *Descartes' error*. Grosset/Putnam. [arLWB, LCC]
- Damasio, A. R. & Damasio, H. (1994) Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In: *Large-scale neuronal theories of the brain*, ed. C. Koch & J. L. Davis. MIT Press. [aLWB]
- D'Andrade, R. (1987) A folk model of the mind. In: *Cultural models in language and thought*, ed. D. Holland & N. Quinn. Cambridge University Press. [aLWB]
- Danesi, M. (1998) The “dimensionality principle” and semiotic analysis. *Sign Systems Studies* 26:42–60. [AT]
- Davies, M. & Stone, T., eds. (1995) *Mental simulation*. Basil Blackwell. [aLWB]
- Davis, M. (1998) Are different parts of the extended amygdala involved in fear versus anxiety? *Biological Psychiatry* 44:1239–47. [rLWB]
- DeJong, G. & Mooney, R. (1986) Explanation-based learning: An alternative view. *Machine Learning* 1:145–76. [aLWB]
- Dennett, D. C. (1969) The nature of images and the introspective trap. In: Dennett, D. C. *Content and consciousness*. Routledge. [aLWB]
- Dennett, D. C. & Kinsbourne, M. (1992) Time and the observer - the where and when of consciousness in the brain. *Behavioral and Brain Sciences* 15:183–247. [aLWB]

- Dent-Read, C. & Szkolozsky, A. (1993) Where do metaphors come from? *Metaphor and Symbolic Activity* 8(3):227–42. [BI]
- DeRenzi, E. & Spinnler, H. (1967) Impaired performance on color tasks in patients with hemispheric lesions. *Cortex* 3:194–217. [aLWB]
- Deschaumes-Molinario, C., Dittmer, A. & Vernet-Maury, E. (1992) Autonomic nervous system response patterns correlate with mental imagery. *Physiology and Behavior* 51:1021–27. [aLWB]
- De Sousa, R. (1987) *The rationality of emotion*. Bradford Books, MIT Press. [LCC]
- Diamond, A., Werker, J. F. & Lalonde, M. A. (1994) Toward understanding commonalities in the development of object search, detour navigation, categorization, and speech perception. In: *Human behavior and the developing brain*, ed. G. Dawson & K. W. Fischer. Guilford. [LM]
- Dietrich, E. & Markman, A. B., eds. (in press) *Cognitive dynamics: Conceptual change in humans and machines*. Erlbaum. [rLWB]
- Donald, M. (1991) *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Harvard University Press. [arLWB, LG, AJW]
- (1993) Précis of *Origins of the modern mind: Three stages in the evolution of culture and cognition*. *Behavioral and Brain Sciences* 16:739–91. [arLWB]
- Dretske, F. (1981) *Knowledge and the flow of information*. MIT Press/Bradford Books. [FA, MA]
- (1988) *Explaining behavior*. MIT Press/Bradford Books. [FA]
- (1995) *Naturalizing the mind*. MIT Press. [arLWB]
- Droz, R. & Rahmy, M. (1974) *Lire Piaget*. Charles Dessart. [FL]
- Edelman, G. M. (1987) *Neural Darwinism: The theory of neuronal group selection*. Basic Books. [DAS]
- (1992) *Bright air, brilliant fire: On the matter of the mind*. Basic Books. [aLWB]
- Edelman, S. (1994) Biological constraints and the representation of structure in vision and language. *Psychology* 5(57). <ftp://ftp.princeton.edu/pub/harnad/Psycology/1994.volume.5/psyc.94.5.57.language-network.3.edelman>. [SE]
- (1998) Representation is representation of similarities. *Behavioral and Brain Sciences* 21(4):449–98. [aLWB]
- Ekman, P. (1992) An argument for basic emotions. *Cognition and Emotion* 6:169–200. [LCC]
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996) *Rethinking innateness: A connectionist perspective on development*. MIT Press. [arLWB]
- Epicurus (1994) *The Epicurus reader: Selected writings and testimonia*, trans. B. Inwood & L. P. Gerson. (Original work published in the 4th century B. C.) Hackett. [aLWB]
- Farah, M. J. (1988) Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review* 95:307–17. [aLWB]
- (1995) The neural bases of mental imagery. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aLWB]
- Farah, M. J. & Feinberg, T. E. (1997) Perception and awareness. In: *Behavioral neurology and neuropsychology*, ed. T. E. Feinberg & M. J. Farah. McGraw-Hill. [aLWB]
- Fauconnier, G. (1985) *Mental spaces: Aspects of meaning construction in natural language*. MIT Press. [aLWB, RWL]
- (1997) *Mappings in thought and language*. Cambridge University Press. [aLWB, GF, RWL]
- Fauconnier, G. & Sweetser, E., eds. (1996) *Spaces, worlds, and grammar*. University of Chicago Press. [GF, RWL]
- Fauconnier, G. & Turner, M. (1998) Conceptual integration networks. *Cognitive Science* 22(2):133–87. [arLWB, GF, RWL]
- Fillmore, C. J. (1985) Frames and the semantics of understanding. *Quaderni di Semantica* 6:222–55. [aLWB]
- Fincher-Kiefer, R. (1998) Visuospatial working memory in mental model construction. *Abstracts of the Psychonomic Society* 3:79. [RAZ]
- Finke, R. A. (1989) *Principles of mental imagery*. MIT Press. [aLWB]
- Flavell, J. (1976) Metacognitive aspects of problem solving. In: *The nature of intelligence*, ed. L. Resnick. Erlbaum. [RO]
- Fodor, J. A. (1975) *The language of thought*. Harvard University Press. [aLWB, DCD]
- (1981a) *Representations: Philosophical essays on the foundations of cognitive science*. MIT Press. [ABM]
- (1981b) The present status of the innateness controversy. In: *Representations*, J. A. Fodor. MIT Press. [BL]
- (1983) *The modularity of mind: An essay on faculty psychology*. Bradford Books, MIT Press. [aLWB]
- (1987) *Psychosemantics: The problem of meaning in the philosophy of mind*. MIT Press. [MA]
- (1990) A theory of content (I and II). In: *A theory of content and other essays*, J. A. Fodor. MIT Press. [MA]
- Fodor, J. A. & LePore, E. (1992) *Holism: A shopper's guide*. Blackwell. [aLWB]
- Fodor, J. A. & McLaughlin, B. (1990) Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition* 35(2):183–204. [ABM]
- Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [FA, arLWB, ABM]
- Fox, N. A. & Bell, M. A. (1993) Frontal functioning in cognitive and emotional behaviors during infancy: Effects of maturation and experience. In: *NATO ASI Series: D. Developmental neurocognition: Speech and face processing in the first year of life*, ed. S. de Schonen, P. Juszczyk, P. McNeilage & J. Morton. Kluwer. [LM]
- Frege, G. (1892/1952) On sense and reference. In: *Philosophical writings of Gottlieb Frege*, ed. P. T. Geach & M. Black. (Original work published in 1892). Blackwell. [aLWB]
- Freksa, C. (1991) Qualitative spatial reasoning. In: *Cognitive and linguistic aspects of geographic space*, ed. D. M. Mark & A. U. Frank. Kluwer. [CF]
- (1997) Spatial and temporal structures in cognitive processes. In: *Foundations of computer science: Potential - theory - cognition*, ed. C. Freksa, M. Jantzen & R. Valk. Springer. [CF]
- Freksa, C. & Barkowsky, T. (1996) On the relation between spatial concepts and geographic objects. In: *Geographic objects with indeterminate boundaries*, ed. P. Burrough & A. Frank. Taylor & Francis. [CF]
- Freyd, J. J. (1987) Dynamic mental representations. *Psychological Review* 94:427–38. [aLWB]
- Frith, C. & Dolan, R. J. (1997) Brain mechanisms associated with top-down processes in perception. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences* 352:1221–30. [aLWB]
- Gabora, L. (1998) Autocatalytic closure in a cognitive system: A tentative scenario for the evolution of culture. *Psychology* 9(67). www.cogsci.soton.ac.uk/psych/newpsy9?9.67 [ftp://ftp.princeton.edu/pub/harnad/Psycology/1998.volume.9/page.98.9.XX.-1.gabora](http://ftp.princeton.edu/pub/harnad/Psycology/1998.volume.9/page.98.9.XX.-1.gabora). [LG]
- (1999) Weaving, bending, patching, mending the fabric of reality: A cognitive science perspective on worldview inconsistency. *Foundations of Science* 3(2). www.vub.ac.be/clea/liane/fabric_of_reality/html. [LG]
- Gainotti, G., Silveri, M. C., Daniele, A. & Giustolisi, L. (1995) Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory* 3:247–64. [aLWB]
- Garner, W. R. (1974) *The processing of information and structure*. Wiley. [aLWB]
- (1978) Aspects of a stimulus: Features, dimensions, and configurations. In: *Cognition and categorization*, ed. E. Rosch & B. B. Lloyd. Erlbaum. [aLWB]
- Gazzaniga, M. S. (1988) Brain modularity: Towards a philosophy of conscious experience. In: *Consciousness in contemporary science*, ed. A. J. Marcel & E. Bisiach. Clarendon Press. [aLWB]
- Gazzaniga, M. S., Ivry, R. B. & Mangun, G. R. (1998) *Cognitive neuroscience: The biology of the mind*. Norton. [aLWB]
- Geach, P. T. (1957) *Mental acts*. Routledge. [aLWB]
- Gelman, S. A., Coley, J. D., Rosengran, K. S., Hartman, E. & Pappas, A. (1998) Beyond labeling: The role of maternal input in the acquisition of richly structured categories. *Monographs of the Society for Research in Child Development* 63, Serial No. 253. [rLWB]
- Gelman, S. A. & Diesendruck, G. (in press) What's in a concept? Context, variability, and psychological essentialism. In: *Theoretical perspectives in the concept of representation*, ed. I. E. Siegel. Erlbaum. [rLWB]
- Gelman, S. A. & Tardiff, T. (1998) A cross-linguistic comparison of generic noun phrases in English and Mandarin. *Cognition* 66:215–48. [rLWB]
- Gentner, D. (1983) Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7:155–70. [LG, ABM]
- Gentner, D. & Markman, A. B. (1997) Structural alignment in analogy and similarity. *American Psychologist* 52:45–56. [rLWB, LG, ABM]
- Gentner, D. & Stevens, A. L., eds. (1983) *Mental models*. Erlbaum. [aLWB]
- Gernsbacher, M. A., Varner, K. R. & Faust, M. E. (1990) Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:430–45. [aLWB]
- Gibbs, R. W., Jr. (1983) Do people always process the literal meanings of indirect requests? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9:524–33. [aLWB]
- (1992) What do idioms really mean? *Journal of Memory and Language* 31:485–506. [RWG]
- (1994) *The poetics of mind: Figurative thought, language, and understanding*. Cambridge University Press. [aLWB, GF, RWG]
- Gibson, J. J. (1979) *The ecological approach to visual perception*. Houghton Mifflin. [aLWB]
- Gilbert, D. T. (1991) How mental systems believe. *American Psychologist* 46:107–19. [rLWB]
- Gilden, D., Blake, R. & Hurst, G. (1995) Neural adaptation of imaginary visual motion. *Cognitive Psychology* 28:1–16. [aLWB]
- Gineste, M.-D., Indurkha, B. & Scart-Lhomme, V. (1997) Mental representations in understanding metaphors. *Notes et Documents Limsi* No. 97–02, LIMSI-CNRS, BP 133, F-91403, Orsay, Cedex, France. [BI]
- Givón, T. (1978) Negation in language: Pragmatics, function, ontology. In: *Syntax and semantics 9: Pragmatics*, ed. P. Cole. Academic Press. [aLWB]

- Glaser, W. R. (1992) Picture naming. *Cognition* 42:61–105. [arLWB]
- Glasgow, J. I. (1993) The imagery debate revisited: A computational perspective. *Computational Intelligence* 9:309–33. [aLWB]
- Glasgow, J. I., Narayanan, H. & Chandrasekaran, B., eds. (1995) *Diagrammatic reasoning: Computational and cognitive perspectives*. MIT Press. [CF]
- Glenberg, A. M. (1997) What memory is for. *Behavioral and Brain Sciences* 20:1–55. [arLWB, AMG]
- Glenberg, A. M., Meyer, M. & Lindem, K. (1987) Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Language* 26:69–83. [aLWB]
- Glenberg, A. M. & Robertson, D. A. (in press) *Indexical understanding of instructions*. Discourse Processes. [AMG]
- Glenberg, A. M., Robertson, D. A., Jansen, J. L. & Johnson, M. C. (1998a) Not propositions. *Journal of Cognitive Systems* (under review). [aLWB]
- Glenberg, A. M., Robertson, D. A. & Members of the Honors Seminar in Cognitive Psychology, 1997–98 (1998b) Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language* (under review). [aLWB, AMG]
- Glenberg, A.M., Schroeder, J. L. & Robertson, D. A. (1998) Averting the gaze disengages the environment and facilitates remembering. *Memory and Cognition* 26:651–58. [rLWB]
- Goldstone, R. (1994) Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General* 123:178–200. [aLWB]
- Goldstone, R. & Barsalou, L. W. (1998) Reuniting cognition and perception: The perceptual bases of rules and similarity. *Cognition* 65:231–62. [aLWB]
- Goodman, N. (1976) *Languages of art*. Hackett. [aLWB, BI]
- (1978) *Ways of worldmaking*. Hackett. [BI]
- Gopnik, A. & Slaughter, V. (1991) Young children's understanding of changes in their mental states. *Child Development* 62:98–110. [RO]
- Gopnik, A. & Wellman, H. M. (1994) The theory theory. In: *Mapping the mind. Domain specificity in cognition and culture*, ed. L. A. Hirschfeld & S. A. Gelman. Cambridge University Press. [RO]
- Gordon, R. M.. (1987) *The structure of emotion: Investigations in cognitive philosophy*. Cambridge University Press. [LCC]
- Gordon, W. J. J. (1961) *Synectics: The development of creative capacity*. Harper & Row. [BI]
- (1965) The metaphorical way of knowing. In: *Education of vision*, ed. G. Kepes. George Braziller. [BI]
- Gould, S. J. (1991) Exaptation: A crucial tool for an evolutionary psychology. *Journal of Social Issues* 47:43–65. [aLWB]
- Grady, J. (1997) Theories are buildings revisited (metaphor). *Cognitive Linguistics* 8:267–90. [RWG]
- Graesser, A. C. & Clark, L. F. (1985) *Structures and procedures of implicit knowledge*. Ablex. [KW-H]
- Grossberg, S. (1987) Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science* 11:23–63. [DAS]
- Hadamard, J. (1949) *The psychology of invention in the mathematical field*. Dover Books. [aLWB]
- Haith, M. M. (1994) Visual expectations as the first step toward the development of future-oriented processes. In: *The John D. and Catherine T. MacArthur Foundation series on mental health and development*, ed. M. M. Haith, J. B. Benson & R. J. Roberts, Jr. University of Chicago Press. [aLWB]
- Hala, S. & Carpendale, J. (1977) All in the mind: Children's understanding of mental life. In: *The development of social cognition*, ed. S. Hala. Psychology Press. [rLWB]
- Half, H. M., Ortony, A. & Anderson, R. C. (1976) A context-sensitive representation of word meanings. *Memory and Cognition* 4:378–83. [aLWB]
- Hampton, J. A. (1981) An investigation of the nature of abstract concepts. *Memory and Cognition* 9:149–56. [KW-H]
- Harnad, S. (1987) Category induction and representation. In: *Categorical perception: The groundwork of cognition*, ed. S. Harnad. Cambridge University Press. [arLWB]
- (1990) The symbol grounding problem. *Physica D* 42:335–46. [arLWB, SE]
- Harré, R. (1970) *Principles of scientific thinking*. University of Chicago Press. [WFB]
- Haugeland, J. (1985) *Artificial intelligence: The very idea of it*. MIT Press. [aLWB]
- (1991) Representational genera. In: *Philosophy and connectionist theory*, ed. W. Ramsey, S. P. Stich & D. E. Rumelhart. Erlbaum. [aLWB]
- Hausman, C. R. (1984) *A discourse on novelty and creation*. SUNY Press. [BI]
- (1989) *Metaphor and art: Interactionism and reference in verbal and nonverbal art*. Cambridge University Press. [BI]
- Heal, J. (1996) Simulation and cognitive penetrability. *Mind and Language* 11:44–67. [aLWB]
- Hebb, D. O. (1949) *The organization of behavior: A neurophysiological approach*. Wiley. [JH, DAS]
- Heit, E. & Barsalou, L. W. (1996) The instantiation principle in natural categories. *Memory* 4:413–41. [aLWB]
- Helmholtz, H. von (1962) *Treatise in physiological optics, vol. III*, trans. J. P. Southall. Dover. (Original work published in 1866). [JH]
- Herskovits, A. (1997) Language, spatial cognition, and vision. In: *Temporal and spatial reasoning*, ed. O. Stock. Kluwer. [aLWB]
- Hespos, S. J. & Rochat, P. (1997) Dynamic mental representation in infancy. *Cognition* 64:153–88. [aLWB]
- Hesse, M. (1967) Models and analogy in science. In: *Encyclopedia of philosophy, vol. 5*, ed. P. Edwards. Macmillan. [WFB]
- Hillis, A. E. & Caramazza, A. (1995) Cognitive and neural mechanisms underlying visual and semantic processing: Implications from "optic aphasia." *Journal of Cognitive Neuroscience* 7:457–78. [aLWB]
- Hochberg, J. (1968) In the mind's eye. In: *Contemporary theory and research in visual perception*, ed. R. N. Haber. Holt, Rinehart & Winston. [JH]
- (1970) Attention, organization, and consciousness. In: *Attention: Contemporary theory and analysis*, ed. D. I. Mostofsky. Appleton-Century-Crofts. [JH]
- (1998) Gestalt theory and its legacy: Organization in eye and brain, in attention and mental representation. In: *Perception and cognition at century's end: Handbook of perception and cognition, 2nd edition*, ed. J. Hochberg. Academic Press. [rLWB, JH]
- Hochberg, J. & Brooks, V. (1996) The perception of motion pictures. In: *Handbook of perception and cognition: Cognitive ecology*, ed. M. Friedman & E. Carterette. Academic Press. [JH]
- Höfding, H. (1891) *Outlines of psychology*. Macmillan. [arLWB]
- Hoffman, J. E. & Subramaniam, B. (1995) Saccadic eye movements and selective visual attention. *Perception and Psychophysics* 57:787–95. [JH]
- Hofstadter, D. (1995) *Fluid concepts and creative analogies*. Basic Books. [GF]
- Houk, J. C., Buckingham, J. T. & Barto, A. G. (1996) Models of cerebellum and motor learning. *Behavioral and Brain Sciences* 19(3):368–83. [NN]
- Hume, D. (1739/1978) *A treatise on human nature*, 2nd edition. Oxford University Press. [aLWB]
- Hummel, J. E. & Holyoak, K. J. (1997) Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* 104:427–66. [rLWB]
- Huttenlocher, J. (1973) Language and thought. In: *Communication, language, and meaning: Psychological perspectives*, ed. G. A. Miller. Basic Books. [aLWB]
- (1976) Language and intelligence. In: *The nature of intelligence*, ed. L. B. Resnick. Erlbaum. [aLWB]
- Huttenlocher, J., Jordan, N. C. & Levine, S. C. (1994) A mental model for early arithmetic. *Journal of Experimental Psychology: General* 123:284–96. [aLWB]
- Huttenlocher, P. R. (1994) Synaptogenesis in human cerebral cortex. In: *Human behavior and the developing brain*, ed. G. Dawson & K. W. Fischer. Guilford. [LM]
- Indurkha, B. (1992) *Metaphor and cognition: An interactionist approach*. Kluwer. [GF, BI]
- (1998) On creation of features and change of representation. *Journal of Japanese Cognitive Science Society* 5(2):43–56. [BI]
- Intraub, H., Gottesman, C. V. & Bills, A. J. (1998) Effects of perceiving and imagining scenes on memory for pictures. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24:186–201. [aLWB]
- Intraub, H. & Hoffman, J. E. (1992) Reading and visual memory: Remembering scenes that were never seen. *American Journal of Psychology* 105:101–14. [aLWB]
- Irwin, D. E. (1996) Integrating information across saccadic eye movements. *Current Directions in Psychological Science* 5:94–100. [JH]
- Jackendoff, R. (1987) On beyond zebra: The relation of linguistic and visual information. *Cognition* 26:89–114. [aLWB]
- Jacoby, L. (1983) Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior* 22:485–508. [aLWB]
- Jacoby, L. L., Kelley, C. M., Brown, J. & Jasechko, J. (1989) Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology* 56:326–38. [aLWB]
- James, W. (1890) *Principles of psychology*. Holt. [RO]
- Jeannerod, M. (1994) The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences* 17:187–245. [aLWB]
- (1995) Mental imagery in the motor context. *Neuropsychologia* 33:1419–32. [aLWB]
- (1997) *The cognitive science of action*. Blackwell. [NN]
- Johnson, M. (1987) *The body in the mind: The bodily basis of meaning, imagination, and reason*. University of Chicago Press. [aLWB, RWL]
- Johnson, M. K. & Raye, C. L. (1981) Reality monitoring. *Psychological Review* 88:67–85. [arLWB]
- Johnson-Laird, P. N. (1983) *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press. [aLWB, LCC, RAZ]
- Jones, S. S. & Smith, L. B. (1993) The place of perception in children's concepts. *Cognitive Development* 8:113–39. [aLWB]

- Jonides, J. & Smith, E. E. (1997) The architecture of working memory. In: *Cognitive neuroscience*, ed. M. D. Rugg. MIT Press. [aLWB]
- Jorion, P. (1990) *Principes des systèmes intelligents*. Collection "Sciences cognitives." Masson. (Summary in English: <http://cogprints.soton.ac.uk/abs/comp/199806039>). [PJM]
- (1999) Le secret de la chambre chinoise. *L'Homme* 150:1–26. [PJM]
- Just, M. A. & Carpenter, P. A. (1987) *The psychology of reading and language comprehension*. Allyn & Bacon. [aLWB]
- Kahnemann, D. & Tversky, A. (1982) The simulation heuristic. In: *Judgement under uncertainty: Heuristics and biases*, ed. D. Kahnemann, P. Slovic & A. Tversky. Cambridge University Press. [aLWB]
- Kanizsa, G. (1979) *Organization in vision: Essays in Gestalt perception*. Praeger Press. [aLWB]
- Kant, E. (1787/1965) *The critique of pure reason*, trans. N. K. Smith. St. Martin's Press. [aLWB]
- Kaplan, S., Sonntag, M. & Chown, E. (1991) Tracing recurrent activity in cognitive elements (TRACE): A model of temporal dynamics in a cell assembly. *Connection Science* 3:179–206. [DAS]
- Kaplan, S., Weaver, M. & French, R. (1990) Active symbols and internal models: Towards a cognitive connectionism. *A I and Society* 4:51–71. [DAS]
- Kauffman, S. A. (1993) *Origins of order*. Oxford University Press. [LG]
- Keil, F. C. (1989) *Concepts, kinds, and cognitive development*. MIT Press. [aLWB]
- Kellman, P. (1993) Kinematic foundations of infant visual perception. In: *Visual perception and cognition in infancy*, ed. C. Granrud. Erlbaum. [LM]
- Kintsch, W. (1974) *The representation of meaning in memory*. Erlbaum. [aLWB]
- Klatzky, R. L., Pellegrino, J. W., McCloskey, B. P. & Doherty, S. (1989) Can you squeeze a tomato? The role of motor representations in semantic sensibility judgments. *Journal of Memory and Language* 28:56–77. [RAZ]
- Koestler, A. (1964) *The act of creation*. (2nd Danube edition, 1976). Hutchinsons of London. [BI]
- Kolers, P. A. & Roediger, H. L., III. (1984) Procedures of mind. *Journal of Verbal Learning and Verbal Behavior* 23:425–49. [aLWB]
- Kosslyn, S. M. (1976) Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval time. *Memory and Cognition* 4:291–97. [aLWB]
- (1980) *Image and mind*. Harvard University Press. [aLWB, DCD]
- (1987) Seeing and imaging in the cerebral hemispheres: A computational approach. *Psychological Review* 94:148–75. [aLWB]
- (1994) *Image and brain*. MIT Press. [aLWB]
- Kosslyn, S. M., Cave, C. B., Provost, D. A. & von Gierke, S. M. (1988) Sequential processes in image generation. *Cognitive Psychology* 20:319–43. [aLWB]
- Kosslyn, S. M., Thompson, W. L., Kim, I. J. & Alpert, N. M. (1995) Topographical representations of mental images in primary visual cortex. *Nature* 378:496–98. [aLWB]
- Kowler, E., Anderson, E., Doshier, B. & Blasser, E. (1995) The role of attention in the programming of saccades. *Vision Research* 35:1897–916. [JH]
- Lakoff, G. (1987) *Women, fire, and dangerous things: What categories reveal about the mind*. University of Chicago Press. [aLWB, CF, RWL]
- (1988) Cognitive semantics. In: *Meaning and mental representations*, ed. U. Eco, M. Santambrogio & P. Violi. Indiana University Press. [aLWB]
- (1990) The invariance hypothesis: Is abstract reason based on image-schemas? *Cognitive Linguistics* 1:39–74. [RWL]
- Lakoff, G. & Johnson, M. (1980) *Metaphors we live by*. University of Chicago Press. [aLWB, RWL]
- (1999) *Philosophy in the flesh*. Basic Books. [CF]
- Lakoff, G. & Turner, M. (1989) *More than cool reason: A field guide to poetic metaphor*. University of Chicago Press. [aLWB, RWL]
- 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* 104:211–40. [aLWB, AMG, TKL]
- Landauer, T. K., Foltz, P. W. & Laham, D. (1998a) An introduction to Latent Semantic Analysis. *Discourse Processes* 25:259–84. [TKL]
- Landauer, T. K., Laham, D. & Foltz, P. W. (1998b) Learning human-like knowledge by singular value decomposition: A progress report. *Advances in Neural Information Processing Systems* 10:45–51. [TKL]
- Lang, P. J. (1979) A bio-informational theory of emotional imagery. *Psychophysiology* 16:495–512. [RAZ]
- Langacker, R. W. (1986) An introduction to cognitive grammar. *Cognitive Science* 10:1–40. [aLWB]
- (1987) *Foundations of cognitive grammar, vol. 1. Theoretical prerequisites*. Stanford University Press. [aLWB, CF, RWL]
- (1990) *Concept, image, and symbol: The cognitive basis of grammar*. Mouton de Gruyter. [RWL]
- (1991) *Foundations of cognitive grammar, vol. 2. Descriptive application*. Stanford University Press. [aLWB, CF, RWL]
- (1993) Universals of construal. *Proceedings of the Annual Meeting of the Annual Meeting of the Berkeley Linguistics Society* 19:447–63. [RWL]
- (1995) Viewing in cognition and grammar. In: *Alternative linguistics: Descriptive and theoretical modes*, ed. P. W. Davis. John Benjamins. [RWL]
- (1997) Constituency, dependency, and conceptual grouping. *Cognitive Linguistics* 8:1–32. [aLWB]
- (1998) Conceptualization, symbolization, and grammar. In: *The new psychology of language: Cognitive and functional approaches to language structure*, ed. M. Tomasello. Erlbaum. [RWL]
- Larbig, W., Montoya, P., Flor, H., Bilow, H., Weller, S. & Birbaumer, N. (1996) Evidence for a change in neural processing in phantom limb pain patients. *Pain* 67:275–83. [SE]
- Lassaline, M. E. & Logan, G. D. (1993) Memory-based automaticity in the discrimination of visual numerosity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19:561–81. [aLWB]
- Lazarus, R. (1991) *Emotion and adaptation*. Oxford University Press. [rLWB, LCC]
- LeDoux, J. E. (1996) *The emotional brain: The mysterious underpinnings of emotional life*. Simon & Schuster. [aLWB, LCC]
- Lehrer, A. (1974) *Semantic fields and lexical structure*. American Elsevier. [aLWB]
- Lehrer, A. & Kittay, E. F., eds. (1992) *Frames, fields, and contrasts: New essays in lexical and semantic organization*. Erlbaum. [aLWB]
- Lehrer, K. (1989) *Thomas Reid*. Routledge. [aLWB]
- Leslie, A. M. (1987) Pretence and representation: The origins of "theory of mind." *Psychological Review* 94:412–26. [RO]
- (1994) ToMM, ToBy, and Agency: Core architecture and domain specificity. In: *Mapping the mind: Domain specificity in cognition and culture*, ed. L. A. Hirschfeld & S. A. Gelman. Cambridge University Press. [RO]
- Leslie, A. M. & German, T. P. (1995) Knowledge and ability in "theory of mind": One-eyed overview of a debate. In: *Mental simulation. Evaluations and applications*, ed. M. Davies & T. Stone. Blackwell. [RO]
- Levin, B. (1995) *English verb classes and alternations*. University of Chicago Press. [aLWB]
- Levine, D. N., Warach, J. & Farah, M. J. (1985) Two visual systems in mental imagery: Dissociation of "What" and "Where" in imagery disorders due to bilateral posterior cerebral lesions. *Neurology* 35:1010–18. [aLWB]
- Libet, B. (1982) Brain stimulation in the study of neuronal functions for conscious sensory experiences. *Human Neurobiology* 1:235–42. [aLWB]
- (1985) Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences* 8:529–66. [aLWB]
- Locke, J. (1690/1959) *An essay concerning human understanding, vols. I and II*. 1st edition. Dover. (Original work published in 1690). [aLWB]
- Logan, G. (1988) Toward an instance theory of automatization. *Psychological Review* 95:492–527. [aLWB]
- (1995) Linguistic and conceptual control of visual spatial attention. *Cognitive Psychology* 28:103–74. [aLWB]
- Logan, G. D. & Etherton, J. L. (1994) What is learned during automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:1022–50. [aLWB]
- Logan, G. D., Taylor, S. W. & Etherton, J. L. (1996) Attention in the acquisition and expression of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:620–38. [aLWB]
- Lotman, Y. M. (1983) Asimmetrija i dialog. [Asymmetry and a dialogue]. *Trudy po Znakovym Sistemam* 16:15–30. [AT]
- (1984) O semiosfere. [On a semiosphere]. *Trudy po Znakovym Sistemam* 17:5–23. [AT]
- (1989) O roli sluchainykh faktorov v literaturnoi evoljutsii. [On the role of chance factors in the evolution of literature]. *Trudy po Znakovym Sistemam* 23:39–48. [AT]
- Lowenthal, F. (1978) Logic of natural language and games at primary school. *Revue de Phonétique Appliquée* 46/47:133–40. [FL]
- (1986) Non-verbal communication devices: Their relevance, their use and the mental processes involved. In: *Pragmatics and education*, ed. F. Lowenthal & F. Vandamme. Plenum Press. [FL]
- (1992) Formalism for clinical observations. In: *Advances in nonverbal communication*, ed. F. Poyatos. John Benjamins. [FL]
- Lowenthal, F. & Saerens, J. (1982) Utilisation de formalismes logiques pour l'examen d'enfants aphasiques. *Acta Neurologica Belgica* 16:275–80. [FL]
- (1986) Evolution of an aphasic child after the introduction of NVCDs. In: *Pragmatics and education*, ed. F. Lowenthal & F. Vandamme. Plenum Press. [FL]
- Luria, A. R. (1969) *Vyshije korkovije funkcii celoveka i ih narushenija pri lokal'nyh porazenijah mozga*. [Higher cortical functions in man and their disturbances in local brain lesions]. Moscow University Press. [AT]
- (1973) *Osnovy neiropsikhologii*. [Foundations of neuropsychology]. Moscow University Press. [AT]
- (1979) *Jazyk i soznanije*. [Language and consciousness]. Moscow University Press. [AT]
- Lyons, W. (1980) *Emotion*. Cambridge University Press. [LCC]

- MacDonald, M. C. & Just, M. A. (1989) Changes in activation level with negation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15:633–42. [RAZ]
- MacWhinney, B. (1998) *The emergence of language*. Erlbaum. [arLWB]
- Magnuson, J. S. & Nusbaum, H. C. (1993) Talker differences and perceptual normalization. *Journal of the Acoustical Society of America* 93:2371. [aLWB]
- Malt, B. C. (1994) Water is not H₂O. *Cognitive Psychology* 27:41–70. [rLWB]
- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M. & Wang, Y. (1999) Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language* 40:230–62. [rLWB]
- Mandelblat, N. (1997) *Creativity and schematicity in grammar and translation: The cognitive mechanisms of blending*. Doctoral dissertation, Cognitive Science, University of California, San Diego. [GF]
- Mandler, G. (1975) *Mind and body: Psychology of emotion and stress*. W. W. Norton. [aLWB, LCC]
- Mandler, J. M. (1992) How to build a baby: II. Conceptual primitives. *Psychological Review* 99:587–604. [aLWB]
- Marcel, A. J. (1983a) Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology* 15:197–237. [aLWB]
- (1983b) Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology* 15:238–300. [aLWB]
- Marcus, G. F. (1998) Can connectionism save constructivism? *Cognition* 66:153–82. [ABM]
- Markman, A. B. (1999) *Knowledge representation*. Erlbaum. [ABM]
- Markman, A. B. & Gentner, D. (1993) Splitting the differences: A structural alignment view. *Journal of Memory and Language* 32:517–35. [KW-H]
- Markman, E. M. (1989) *Categorization and naming in children*. MIT Press. [arLWB]
- Markman, K. D., Gavinski, I., Sherman, S. J. & McMullen, M. N. (1993) The mental simulation of better and worse possible worlds. *Journal of Experimental Social Psychology* 29:87–109. [aLWB]
- Marr, D. (1982) *Vision*. W. H. Freeman. [ABM]
- Marschark, M., Katz, A. & Paivio, A. (1983) Dimensions of metaphors. *Journal of Psycholinguistic Research* 12:17–40. [BI]
- Marslen-Wilson, W. D. & Tyler, L. K. (1980) The temporal structure of spoken language understanding. *Cognition* 8:1–71. [aLWB]
- Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L. & Ungerleider, L. G. (1995) Discrete cortical regions associated with knowledge of color and knowledge of action. *Science* 270:102–105. [aLWB]
- Martin, A., Wiggs, C. L., Ungerleider, L. G. & Haxby, J. V. (1996) Neural correlates of category-specific knowledge. *Nature* 379:649–52. [aLWB]
- Martin, M. & Jones, G. V. (1998) Generalizing everyday memory: Signs and handedness. *Memory and Cognition* 26:193–200. [aLWB]
- Mauro, A. (1990) *Observation de deux enfants polyhandicapés à l'aide de manipulations concrètes de systèmes formels*. Post-graduat en sciences psychopédagogiques, Université de Mons-Hainaut. [FL]
- McClelland, J. L., Rumelhart, D. E. & the PDP Research Group (1986) *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models*. MIT Press. [aLWB]
- McCloskey, M. (1983) Naïve theories of motion. In: *Mental models*, ed. D. Gentner & A. L. Stevens. Erlbaum. [aLWB]
- McCune, L. (1995) A normative study of representational play at the transition to language. *Developmental Psychology* 31:198–206. [LM]
- McCune-Nicolich, L. (1981) The cognitive basis of relational words. *Journal of Child Language* 8:15–36. [LM]
- McDermott, D. (1987) A critique of pure reason. *Computational Intelligence* 3:151–60. [aLWB]
- McLaughlin, B. (1997) Classical constituents in Smolensky's ICS architecture. In: *Structures and norms in science*, ed. M. L. Dalla Chiara, K. Doets, D. Mundici & J. van Benthem. Kluwer Academic Publishers. [ABM]
- McNamara, T. P. & Sternberg, R. J. (1983) Mental models of word meaning. *Journal of Verbal Learning and Verbal Behavior* 22:449–74. [KW-H]
- Medin, D. L. & Ross, B. H. (1989) The specific character of abstract thought: Categorization, problem-solving, and induction. In: *Advances in the psychology of human intelligence, vol. 5*, ed. R. J. Sternberg. Erlbaum. [aLWB]
- Medin, D. L. & Schaffer, M. (1978) A context theory of classification learning. *Psychological Review* 85:207–38. [aLWB]
- Melara, R. D. & Marks, L. E. (1990) Dimensional interactions in language processing: Investigating directions and levels of crosstalk. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:539–54. [aLWB]
- Melzack, R., Israel, R., Lacroix, R. & Schultz, G. (1997) Phantom limbs in people with congenital limb deficiency or amputation in early childhood. *Brain* 120:1603–20. [SE]
- Miikkulainen, R. (1993) *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. MIT Press. [SE]
- Mill, J. S. (1965) On the permanent possibility of sensation. An examination of Sir William Hamilton's philosophy. In: *A source book in the history of psychology*, ed. R. J. Herrnstein & E. G. Boring. Harvard University Press. [JH]
- Miller, G. A. & Johnson-Laird, P. N. (1976) *Language and perception*. Harvard University Press. [aLWB]
- Millikan, R. G. (1998) A common structure for concepts of individuals, stuffs, and real kinds: More Mama, more milk, and more mouse. *Behavioral and Brain Sciences* 21:55–100. [arLWB]
- Mills, D. L., Coffey-Corina, S. & Neville, H. J. (1994) Variability in cerebral organization during primary language acquisition. In: *Human behavior and the developing brain*, ed. G. Dawson & K. W. Fischer. Guilford Press. [LM]
- Milner, A. D. & Goodale, M. A. (1995) *The visual brain in action*. Oxford University Press. [aLWB]
- Miner, A. & Reder, L. M. (1994) A new look at feeling-of-knowing: Its metacognitive role in regulating question answering. In: *Metacognition: Knowing about knowing*, ed. J. Metcalfe & A. P. Shimamura. Bradford Books/MIT Press. [RO]
- Minsky, M. L. (1977) A framework for representing knowledge. In: *The psychology of computer vision*, ed. P. H. Winston. McGraw-Hill. [aLWB]
- (1985) *The society of mind*. Simon & Schuster. [aLWB]
- Mitchell, R. W. (1994) The evolution of primate cognition: Simulation, self-knowledge, and knowledge of other minds. In: *Hominid culture in primate perspective*, ed. D. Quiatt & J. Itani. University of Colorado Press. [RWM]
- Mitchell, T. M., Keller, R. M. & Kedar-Cabelli, S. T. (1986) Explanation-based generalization: A unifying view. *Machine Learning* 1:47–80. [aLWB]
- Mitchell, W. T. (1990) *The logic of architecture: Design, computation, and cognition*. MIT Press. [aLWB]
- Montague, R. (1973) The proper treatment of quantification in ordinary language. In: *Approaches to natural language*, ed. J. Hintikka. Reidel. [JRH]
- Morris, C. D., Bransford, J. D. & Franks, J. J. (1977) Levels of processing versus test-appropriate strategies. *Journal of Verbal Learning and Verbal Behavior* 16:519–33. [aLWB]
- Morrow, D. G. & Clark, H. H. (1988) Interpreting words in spatial descriptions. *Language and Cognitive Processes* 3:275–91. [RAZ]
- Morrow, D. G., Greenspan, S. L. & Bower, G. H. (1987) Accessibility and situation models in narrative comprehension. *Journal of Memory and Language* 26:165–87. [aLWB]
- Muhlnickel, W., Elbert, T., Taub, E. & Flor, H. (1998) Reorganization of auditory cortex in tinnitus. *Proceedings of the National Academy of Science* 95:10340–43. [SE]
- Murphy, G. L. (1996) On metaphoric representation. *Cognition* 60:173–204. [aLWB]
- Murphy, G. L. & Medin, D. L. (1985) The role of theories in conceptual coherence. *Psychological Review* 92:289–316. [aLWB]
- Narayanan, S. (1997) *KARMA: Knowledge-based active representations for metaphor and aspect*. Ph. D. dissertation, Computer Science Division, University of California, Berkeley. [GF]
- Neisser, U. (1967) *Cognitive psychology*. Prentice-Hall. [arLWB]
- (1976) *Cognition and reality*. Freeman. [aLWB]
- Neisser, U. & Becklen, R. (1975) Selective looking: Attending to visually specified events. *Cognitive Psychology* 7:480–94. [JH]
- Nelson, D. L., Walling, J. R. & McEvoy, C. L. (1979) Doubts about depth. *Journal of Experimental Psychology: Human Learning and Memory* 5:24–44. [aLWB]
- Nelson, K. (1996) *Language in cognitive development: Emergence of the mediated mind*. Cambridge University Press. [arLWB]
- Nelson, T. O. & Narens, L. (1990) Metamemory: A theoretical framework and new findings. In: *The psychology of learning and motivation, vol. 26*, ed. G. Bower. Academic Press. [RO]
- Nersessian, N. J. (1992) In the theoretician's laboratory: Thought experimenting as mental modeling. *Proceedings of the Philosophical Association of America* 2:291–301. [aLWB]
- Newell, A. (1990) *Unified theories of cognition*. Harvard University Press. [aLWB]
- Newell, A. & Simon, H. A. (1972) *Human problem solving*. Prentice-Hall. [aLWB]
- Newton, N. (1996) *Foundations of understanding*. John Benjamins. [arLWB, AMG]
- Nisbett, R. & Wilson, T. (1977) Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84:231–59. [aLWB]
- Norman, D. A. (1976) *Memory and attention: An introduction to human information processing*. Wiley. [aLWB]
- Norman, D. A., Rumelhart, D. E. & the LNR Research Group (1975) *Explorations in cognition*. Freeman. [aLWB]
- Nosofsky, R. M. (1984) Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10:104–14. [aLWB]
- Nueckles, M. & Jantezko, D. (1997) The role of semantic similarity in the

- comprehension of metaphor. In: *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, ed. M. Shafto & P. Langley. Erlbaum. [BI]
- Nusbaum, H. C. & Morin, T. M. (1992) *Paying attention to differences among talkers*. In: *Speech perception, production, and linguistic structure*, ed. Y. Tohkura, Y. Sagisaka & E. Vatikiotis-Bateson. OHM Publishing. [aLWB]
- Oaksford, M. & Chater, N. (1994) A rational analysis of the selection task as optimal data selection. *Psychological Review* 101:608–31. [aLWB]
- Ohlsson, S. (1993) Abstract schemas. *Educational Psychologist* 28:51–66. [SO]
- Ohlsson, S. & Lehtinen, E. (1997) The role of abstraction in the assembly of complex ideas. *International Journal of Educational Research* 27:37–48. [SO]
- O'Regan, J. K. (1992) Solving the real mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology* 46:461–88. [SE]
- Paivio, A. (1971/1979) *Imagery and verbal processes*. 1971 edition, Holt, Rinehart & Winston. [aLWB] 1979 edition, Erlbaum. [BI]
- (1986) *Mental representations: A dual coding approach*. Oxford University Press. [aLWB]
- Palm, G. (1982) *Neural assemblies: An alternative approach to artificial intelligence*. Springer-Verlag. [DAS]
- Palmer, S. E. (1978) Fundamental aspects of cognitive representation. In: *Cognition and categorization*, ed. E. Rosch & B. Lloyd. Erlbaum. [CF]
- Pani, J. R. (1996) Mental imagery as the adaptationist views it. *Consciousness and Cognition* 5:288–326. [rLWB, WFB]
- Panksepp, J. (1982) Toward a general psychobiological theory of emotions. *Behavioral and Brain Sciences* 5(3):407–69. [LCC]
- (1998) *Affective neuroscience*. Oxford University Press. [LCC]
- Papert, S. (1980) *Mindstorms: Children, computers, and powerful ideas*. Basic Books. [FL]
- Parsons, L. M. (1987a) Imagined spatial transformations of one's body. *Journal of Experimental Psychology: General* 116:172–91. [aLWB]
- (1987b) Imagined spatial transformations of one's hands and feet. *Cognitive Psychology* 19:178–241. [aLWB]
- Peacocke, C. (1992) *A study of concepts*. MIT Press. [aLWB]
- Pennington, N. & Hastie, R. (1992) Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology* 62:189–206. [aLWB]
- Pessoa, L., Thompson, E. & Noë, A. (1998) Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences* 21(6):723–802. [aLWB]
- Peterson, M. A. & Gibson, B. S. (1993) Shape recognition inputs to figure-ground organization in three-dimensional displays. *Cognitive Psychology* 3:383–429. [aLWB]
- Piaget, J. (1954) *The construction of reality in the child*. Basic Books. [LM]
- (1945/1962) *Play, dreams, and imitation in childhood*. W. W. Norton. [LM, RWM]
- (1947/1972) *The psychology of intelligence*. Littlefield, Adams & Co. [RWM]
- Pollack, J. (1990) Recursive distributed representations. *Artificial Intelligence* 46:77–105. [rLWB]
- Posner, M. I. (1995) Attention in cognitive neuroscience: An overview. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aLWB]
- Potter, M. C. & Faulconer, B. A. (1975) Time to understand pictures and words. *Nature* 253:437–38. [aLWB]
- Potter, M. C., Kroll, J. F., Yachzel, B., Carpenter, E. & Sherman, J. (1986) Pictures in sentences: Understanding without words. *Journal of Experimental Psychology: General* 115:281–94. [aLWB]
- Potter, M. C., Valian, V. V. & Faulconer, B. A. (1977) Representation of a sentence and its pragmatic implications: Verbal, imagistic, or abstract? *Journal of Verbal Learning and Verbal Behavior* 16:1–12. [aLWB]
- Premack, D. & Woodruff, G. (1978) Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 4:515–26. [RO]
- Price, H. H. (1953) *Thinking and experience*. Hutchinson's Universal Library. [aLWB]
- Prinz, J. J. (1997) *Perceptual cognition: An essay on the semantics of thought*. Doctoral dissertation, University of Chicago. [aLWB]
- Prinz, J. J. & Barsalou, L. W. (in press a) Acquisition and productivity in perceptual symbol systems: An account of mundane creativity. In: *Creativity and computation*, ed. D. M. Peterson & T. Dartnall. MIT Press. [aLWB]
- (in press b) Steering a course for embodied representation. In: *Cognitive dynamics: Conceptual change in humans and machines*, ed. E. Dietrich. Erlbaum.
- Pulvermüller, F. (1999) Words in the brain's language. *Behavioral and Brain Sciences* 22(2):253–336. [aLWB, DAS]
- Putnam, H. (1960) Minds and machines. In: *Dimensions of mind*, ed. S. Hook. Collier Books. [aLWB]
- Pylyshyn, Z. W. (1973) What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80:1–24. [aLWB]
- (1978) Imagery and artificial intelligence. In: *Perception and cognition: Issues in the foundations of psychology*, ed. C. W. Savage. University of Minnesota Press. [aLWB]
- (1981) The imagery debate: Analogue media versus tacit knowledge. *Psychological Review* 88:16–45. [aLWB]
- (1984) *Computation and cognition*. MIT Press. [aLWB]
- (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22(3). [aLWB]
- Quine, W. V. O. (1953) Two dogmas of empiricism. In: *From a logical point of view*. Harvard University Press. [aLWB]
- (1960) *Word and object*. MIT Press. [KW-H]
- Quinn, P. C. & Eimas, P. D. (1996) Perceptual organization and categorization in young infants. In: *Advances in infancy research, vol. 10*, ed. C. Rovee-Collier & L. P. Lipsitt. Ablex. [BL]
- Ramachandran, V. S. & Hirstein, W. (1998) The perception of phantom limbs. The D. O. Hebb lecture. *Brain* 121:1603–30. [SE]
- Ramachandran, V. S. & Rogers-Ramachandran, D. (1996) Synaesthesia in phantom limbs induced with mirrors. *Proceedings of the Royal Society of London B* 263(1369):377–86. [SE]
- 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* 18:435–52. [RO]
- Reed, C. L. & Vinson, N. G. (1996) Conceptual effects on representational momentum. *Journal of Experimental Psychology: Human Perception and Performance* 22:839–50. [aLWB]
- Regier, T. (1996) *The human semantic potential: Spatial language and constrained connectionism*. MIT Press. [aLWB]
- Reid, T. (1785/1969) *Essays on the intellectual powers of man*. MIT Press. (Original work published in 1785). [aLWB]
- (1764/1970) *An inquiry into the human mind on the principles of common sense*. University of Chicago Press. (Original work published in 1764). [aLWB]
- Rinck, M., Hahnel, A., Bower, G. H. & Glowalla, U. (1997) The metrics of spatial situation models. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:622–37. [aLWB]
- Rips, L. J. (1989) Similarity, typicality, and categorization. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony. Cambridge University Press. [aLWB]
- Roediger, H. L., III. & McDermott, K. B. (1993) Implicit memory in normal human subjects. In: *Handbook of neuropsychology, vol. 8*, ed. F. Boller & J. Grafman. Elsevier. [aLWB]
- Romney, A. K., Weller, S. C. & Batchelder, W. H. (1986) Culture as consensus: A theory of culture and informant accuracy. *American Anthropologist* 88:318–38. [aLWB]
- Rosch, E. & Mervis, C. B. (1975) Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7:573–605. [rLWB]
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M. & Boyes-Braem, P. (1976) Basic objects in natural categories. *Cognitive Psychology* 8:382–439. [aLWB]
- Rosenstein, M. & Cohen, P. R. (1998) Concepts from time series. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. AAAI Press. [aLWB]
- Rösler, F., Heil, M. & Hennighausen, E. (1995) Distinct cortical activation patterns during long-term memory retrieval of verbal, spatial, and color information. *Journal of Cognitive Neuroscience* 7:51–65. [aLWB]
- Ross, B. H. (1996) Category learning as problem solving. In: *The psychology of learning and motivation: Advances in research and theory, vol. 35*, ed. D. L. Medin. Academic Press. [aLWB]
- Ross, B. H., Perkins, S. J. & Tenpenny, P. L. (1990) Reminding-based category learning. *Cognitive Psychology* 22:460–92. [aLWB]
- Roth, J. D. & Kosslyn, S. M. (1988) Construction of the third dimension in mental imagery. *Cognitive Psychology* 20:344–61. [aLWB]
- Rumelhart, D. E., McClelland, J. L. & the PDP Research Group (1986) *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. Foundations*. MIT Press. [aLWB]
- Rumelhart, D. E. & Norman, D. A. (1988) Representation in memory. In: *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition*, ed. R. C. Atkinson, R. J. Herrnstein, G. Lindzey & R. D. Luce. Wiley. [aLWB]
- Rumelhart, D. E. & Ortony, A. (1978) The representation of knowledge in memory. In: *Schooling and the acquisition of knowledge*, ed. R. C. Anderson, R. J. Spiro & W. E. Montague. Erlbaum. [aLWB]
- Rushworth, M. F. S. & Owen, A. M. (1998) The functional organisation of the lateral frontal cortex: Conjecture or conjuncture in the electrophysiology literature? *Trends in Cognitive Sciences* 2:46–53. [aLWB]
- Russell, B. (1919a) *Introduction to mathematical philosophy*. George Allen & Unwin. [aLWB]
- (1919b) On propositions: What they are and how they mean. *Aristotelian Society Supplementary* 2:1–43. (Reprinted in: *The collected papers of Bertrand*

- Russell, col. 8: *The philosophy of logical atomism and other essays, 1914–19*, ed. J. G. Slater. George Allen & Unwin. [aLWB]
- Ryle, G. (1949) *The concept of mind*. Barnes and Noble. [aLWB]
- Sachs, J. D. S. (1967) Recognition memory for syntactic and semantic aspects of connected discourse. *Perception and Psychophysics* 2:437–42. [aLWB]
- (1974) Memory in reading and listening to discourse. *Memory and Cognition* 2:95–100. [aLWB]
- Samuel, A. G. (1981) Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General* 110:474–94. [aLWB]
- (1987) Lexical uniqueness effects on phonemic restoration. *Journal of Memory and Language* 26:36–56. [aLWB]
- (1997) Lexical activation produces potent phonemic percepts. *Cognitive Psychology* 32:97–127. [aLWB]
- Sartre, J. (1948) *The psychology of imagination*. Philosophical Library. [LM]
- Scaife, M. & Rogers, Y. (1996) External cognition: How do graphical representations work? *International Journal of Human-Computer Studies* 45:185–213. [CF]
- Schacter, D. L. (1995) Implicit memory: A new frontier for neuroscience. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aLWB]
- Schacter, D. L., McAndrews, M. P. & Moscovitch, M. (1988) Access to consciousness: Dissociations between implicit and explicit knowledge in neuropsychological syndromes. In: *Thought without language*, ed. L. Weiskrantz. Oxford University Press. [aLWB]
- Schank, R. C. (1982) *Dynamic memory*. Cambridge University Press. [ABM]
- Schank, R. C. & Abelson, R. P. (1977) *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Erlbaum. [aLWB]
- Schneider, W. & Shiffrin, R. M. (1977) Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review* 84:1–66. [aLWB]
- Schön, D. A. (1963) *Displacement of concepts*. Humanities Press. [BI]
- (1979) Generative metaphor: A perspective on problem-setting in social policy. In: *Metaphor and thought*, ed. A. Ortony. Cambridge University Press. [BI]
- Schunn, C. D., Reder, L. M., Nhouyvanisvong, A., Richards, D. R. & Strollofino, P. J. (1997) To calculate or not calculate: A source activation confusion (SAC) model of problem-familiarity's role in strategy selection. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:1–27. [RO]
- Schwab, E. C. (1981) *Auditory and phonetic processing for tone analogs of speech*. Doctoral dissertation, State University of New York at Buffalo. [aLWB]
- Schwanenflugel, P. J. (1991) Why are abstract concepts hard to understand? In: *The psychology of word meaning*, ed. P. J. Schwanenflugel. Erlbaum. [rLWB, KW-H]
- Schwanenflugel, P. J., Fabricius, W. V., Noyes, C. R., Bigler, K. D. & Alexander, J. M. (1994) The organization of mental verbs and folk theories of knowing. *Journal of Memory and Language* 33:376–95. [aLWB]
- Schwanenflugel, P. J., Hamishfeiger, K. K. & Stowe, R. W. (1988) Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language* 27:499–520. [rLWB]
- Schwanenflugel, P. J., Martin, M. & Takahashi, T. (in press) The organization of verbs of knowing: Evidence for cultural commonality and variation in theory of mind. *Memory and Cognition*. [aLWB]
- Schwanenflugel, P. J. & Shoben, E. J. (1983) Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9:82–102. [rLWB, KW-H]
- Schwartz, R. (1981) Imagery - there's more to images than meets the eye. In: *Imagery*, ed. N. Block. MIT Press. [aLWB]
- Schyns, P. G., Goldstone, R. L. & Thibaut, J. P. (1998) The development of features in object concepts. *Behavioral and Brain Sciences* 21:1–54. [aLWB]
- Searle, J. R. (1980) Minds, brains, and programs. *Behavioral and Brain Sciences* 3:417–24. [aLWB]
- (1983) *Intentionality*. Cambridge University Press. [FA]
- Seifert, L. S. (1997) Activating representations in permanent memory: Different benefits for pictures and words. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:1106–21. [aLWB]
- Shannon, B. (1987) The non-abstractness of mental representations. *New Ideas in Psychology* 5:117–26. [aLWB]
- Shastri, L. & Ajanagadde, V. (1993) From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal asynchrony. *Behavioral and Brain Sciences* 16:417–94. [rLWB, ABM]
- Shastri, L. & Grannes, D. (1996) A connectionist treatment of negation and inconsistency. *Proceedings of the 18th Annual Conference of the Cognitive Science Society, University of California at San Diego*. Erlbaum. [GF]
- Shaver, P., Schwartz, J., Kirson, D. & O'Connor, C. (1987) Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology* 52:1061–86. [aLWB]
- Shepard, R. N. & Cooper, L. A. (1982) *Mental images and their transformations*. Cambridge University Press. [aLWB]
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171:701–703. [aLWB]
- Shevell, S. K. & He, J. C. (1995) Interocular difference in Rayleigh matches of color normals. In: *Colour vision deficiencies XII*, ed. B. Dunn. Kluwer Academic Press. [aLWB]
- Shiffrin, M. & Freyd, J. J. (1990) Apparent motion of the human body. *Psychological Science* 4:257–64. [aLWB]
- (1993) Timing and apparent motion path choice with human body photographs. *Psychological Science* 6:379–84. [aLWB]
- Shiffrin, R. M. (1988) Attention. In: *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition*, ed. R. C. Atkinson, R. J. Herrnstein, G. Lindzey & R. D. Luce. Wiley. [aLWB]
- Shiffrin, R. M. & Schneider, W. (1977) Controlled and automatic information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review* 84:127–90. [aLWB]
- Smith, E. E. & Medin, D. L. (1981) *Categories and concepts*. Harvard University Press. [aLWB]
- Smith, L. B. & Heise, D. (1992) Perceptual similarity and conceptual structure. In: *Advances in psychology - percepts, concepts, and categories: The representation and processing of information*, ed. B. Burns. Elsevier. [aLWB]
- Smith, L. B., Jones, S. S. & Landau, B. (1992) Count nouns, adjectives, and perceptual properties in children's novel word interpretation. *Developmental Psychology* 28:273–86. [aLWB]
- Smith, V. C., Pokorny, J. & Starr, S. J. (1976) Variability of color mixture data - 1. Interobserver variability in the unit coordinates. *Vision Research* 16:1087–94. [aLWB]
- Smolensky, P. (1990) Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence* 46:159–216. [rLWB, ABM]
- (1995) Reply: Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In: *Connectionism: Debates on psychological explanation*, ed. C. Macdonald & G. Macdonald. Blackwell. [ABM]
- Snippe, H. P. & Koenderink, J. J. (1992) Discrimination thresholds for channel-coded systems. *Biological Cybernetics* 66:543–51. [SE]
- Snodgrass, J. G. (1984) Concepts and their surface representations. *Journal of Verbal Learning and Verbal Behavior* 23:3–22. [aLWB]
- Snodgrass, J. G. & Vanderwart, M. (1980) A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory* 6:174–215. [aLWB]
- Solomon, K. O. (1997) *The spontaneous use of perceptual representations during conceptual processing*. Doctoral dissertation, University of Chicago. [aLWB]
- Solomon, K. O. & Barsalou, L. W. (1999a) Grounding property verification in perceptual simulation. (Under review). [aLWB]
- (1999b) Representing properties locally: Evidence for the perceptual grounding of concepts. (Under review). [aLWB]
- Solomon, R. (1976) *The passions: The myth and nature of human emotion*. Anchor Press, Doubleday. [LCC]
- Spalding, T. L. & Ross, B. H. (1994) Comparison-based learning: Effects of comparing instances during category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:1251–63. [aLWB]
- Spelke, E. S., Breinlinger, K., Macomber, J. & Jacobson, K. (1992) Origins of knowledge. *Psychological Review* 99:605–32. [aLWB]
- Sperling, G. (1960) The information available in brief visual presentations. *Psychological Monographs* 74, Whole no. 498. [JH]
- Squire, L. R., Knowlton, B. & Musen, G. (1993) The structure and organization of memory. *Annual Review of Psychology* 44:453–95. [aLWB]
- Stein, N. L. & Levine, L. J. (1990) Making sense out of emotion: The representation and use of goal-structured knowledge. In: *Psychological and biological approaches to emotion*, ed. N. L. Stein, B. Leventhal & T. Trabasso. Erlbaum. [aLWB]
- Sweetser, E. (1990) *From etymology to pragmatics: Metaphorical and cultural aspects of semantics*. Cambridge University Press. [aLWB]
- Talmy, L. (1983) How language structures space. In: *Spatial orientation: Theory, research, and application*, ed. H. Pick & L. Acredelo. Plenum Press. [aLWB]
- (1985) Lexicalization patterns: Semantic structure in lexical forms. In: *Language topology and syntactic description, vol. III: Grammatical categories and the lexicon*, ed. T. Shopen. Cambridge University Press. [LM]
- (1996) Fictive motion in language and “ception.” In: *Language and space*, ed. P. Bloom, M. Peterson, L. Nadel & M. Garrett. MIT Press. [CF, RWL]
- (1988) Force dynamics in language and cognition. *Cognitive Science* 12:49–100. [aLWB]

- Tarr, M. J. & Pinker, S. (1989) Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology* 21:233–82. [aLWB]
- Tetlock, P. & Belkin, A., eds. (1996) *Counterfactual thought experiments in world politics: Logical, methodological, and psychological perspectives*. Princeton University Press. [GF]
- Thagard, P. (1992) *Conceptual revolutions*. Princeton University Press. [aLWB]
- The Neural Theory of Language Project. (n.d.) <http://www.icsi.berkeley.edu/NTL/>. [GF]
- Theios, J. & Amrhein, P. C. (1989) Theoretical analysis of the cognitive processing of lexical and pictorial stimuli: Reading, naming, and visual and conceptual comparisons. *Psychological Review* 96:5–24. [aLWB]
- Tomasello, M. (1992) *First verbs: A case study of early grammatical development*. Cambridge University Press. [aLWB]
- Tomasello, M., ed. (1998) *The new psychology of language: Cognitive and functional approaches to language structure*. Erlbaum. [GF]
- Tomasello, M. & Call, J. (1997) *Primate cognition*. Oxford University Press. [rLWB]
- Tomasello, M., Kruger, A. & Ratner, H. (1993) Cultural learning. *Behavioral and Brain Sciences* 16:495–552. [aLWB]
- Toomela, A. (1996) How culture transforms mind: A process of internalization. *Culture and Psychology* 2:285–305. [AT]
- Tootell, R. B. H., Silverman, M. S., Switkes, E. & De Valois, R. L. (1982) Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science* 218:902–904. [aLWB]
- Tourangeau, R. & Ripe, L. (1991) Understanding and appreciating metaphors. *Cognition* 11:203–44. [BI]
- Trabasso, T. R. & Bower, G. H. (1968) *Attention in learning: Theory and research*. Wiley. [aLWB]
- Trehub, A. (1977) Neuronal models for cognitive processes: Networks for learning, perception, and imagination. *Journal of Theoretical Biology* 65:141–69. [rLWB]
- (1991) *The cognitive brain*. MIT Press. [rLWB]
- Treisman, A. M. (1969) Strategies and models of selective attention. *Psychological Review* 76:282–99. [aLWB]
- (1993) The perception of features and objects. In: *Attention: Selection, awareness, and control: A tribute to Donald Broadbent*, ed. A. Baddeley & L. Weiskrantz. Clarendon Press. [aLWB]
- Treisman, A. M. & Geffen, G. (1967) Selective attention: Perception or response? *Quarterly Journal of Experimental Psychology* 19:1–17. [JH]
- Turner, M. (1987) *Death is the mother of beauty: Mind, metaphor, criticism*. University of Chicago Press. [RWL]
- (1996) *The literary mind*. Oxford University Press. [aLWB]
- Turner, M. & Fauconnier, G. (1995) Conceptual integration and formal expression. *Metaphor and Symbolic Activity* 10(3):183–204. [BI]
- (1998) Conceptual integration in counterfactuals. In: *Discourse and cognition*, ed. J.-P. Koenig. CSLI Publications. [GF]
- Tversky, B. & Hemenway, K. (1985) Objects, parts, and categories. *Journal of Experimental Psychology: General* 113:169–93. [aLWB]
- Tye, M. (1991) *The imagery debate*. MIT Press. [aLWB]
- (1993) Image indeterminacy. In: *Spatial representation*, ed. N. Eilen, R. McCarthy & B. Brewer. Blackwell. [aLWB]
- Tzourio, N. & Mellet, E. (1998) Imagerie des fonctions cognitives. *Direction des Sciences du Vivant, lettre d'information scientifique* 16. URL: <http://www-dsv.ceva.fr/dsv/ceabio/16/tzourio16—fr.htm>. [FL]
- Ullman, S. (1996) *High-level vision*. MIT Press. [ABM]
- Ungerleider, L. G. & Mishkin, M. (1982) Two cortical visual systems. In: *Analysis of visual behavior*, ed. D. J. Ingle, M. A. Goodale & R. W. J. Mansfield. MIT Press. [aLWB]
- Van Dijk, T. A. & Kintsch, W. (1983) *Strategies of discourse comprehension*. Academic. [aLWB]
- Van Essen, D. C. (1985) Functional organization of primate visual cortex. In: *Cerebral cortex*, ed. A. Peters & E. G. Jones. Plenum. [aLWB]
- Van Gelder, T. (1990) Compositionality: A connectionist variation on a classical theme. *Cognitive Science* 14:355–84. [rLWB]
- Vaughn, H. G. & Kurtzberg, D. (1992) Electrophysiologic indices of human brain maturation and cognitive development. In: *Developmental behavioral neurosciences: The Minnesota Symposium, vol. 24*, ed. M. R. Gunnar & C. A. Nelson. Erlbaum. [LM]
- Velmans, M. (1991) Is human information processing conscious? *Behavioral and Brain Sciences* 14:651–726. [aLWB]
- Von Eckardt, B. & Potter, M. C. (1985) Clauses and the semantic representation of words. *Memory and Cognition* 13:371–76. [aLWB]
- Vosniadou, S. & Brewer, W. F. (1992) Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology* 24:535–85. [WFB]
- (1994) Mental models of the day/night cycle. *Cognitive Science* 18:123–83. [WFB]
- Vygotsky, L. S. (1925/1982) Soznanije kak problema psikhologii povedenija. [Consciousness as a problem for a psychology of behaviour]. In: *Sobranije sochinenij, Tom 1 [Collected works, vol. 1]*, ed. L. S. Vygotsky. Pedagogika. [AT]
- (1934/1996) *Myshlenije i rech. [Thinking and speech]*. Labirint. [AT]
- (1962) *Thought and language*. MIT Press. [FL]
- Warren, R. M. (1970) Perceptual restoration of missing speech sounds. *Science* 167:392–93. [aLWB]
- Warrington, E. K. & Shallice, T. (1984) Category-specific semantic impairments. *Brain* 107:829–54. [aLWB]
- Watson, J. B. (1913) Psychology as the behaviorist sees it. *Psychological Review* 20:158–77. [aLWB]
- Wattenmaker, W. D. & Shoben, E. J. (1987) Context and the recallability of concrete and abstract sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13:140–50. [rLWB]
- Weiskrantz, L. (1986) *Blindsight: A case study and implications*. Oxford University Press. [aLWB]
- Weisstein, N. & Harris, C. S. (1974) Visual detection of line segments: An object-superiority effect. *Science* 186:752–55. [aLWB]
- (1980) Masking and unmasking of distributed representations in the visual system. In: *Visual coding and adaptability*, ed. C. S. Harris. Erlbaum. [aLWB]
- Weisstein, N., Montalvo, F. S. & Ozog, G. (1972) Differential adaptation to gratings blocked by cubes and gratings blocked by hexagons: A test of the neural symbolic activity hypothesis. *Psychonomic Science* 27:89–91. [aLWB]
- Weisstein, N. & Wong, E. (1986) Figure-ground organization and the spatial and temporal responses of the visual system. In: *Pattern recognition by humans and machines: Vol. 2. Visual perception*, ed. E. C. Schwab & H. C. Nusbaum. Academic Press. [aLWB]
- Wellman, H. & Bartsch, K. (1988) Young children's reasoning about beliefs. *Cognition* 30:239–77. [RO]
- Wellman, H. M. & Gelman, S. A. (1988) Children's understanding of the non-obvious. In: *Advances in the psychology of human intelligence*, ed. R. Sternberg. Erlbaum. [aLWB]
- Wells, A. J. (1998) Turing's analysis of computation and theories of cognitive architecture. *Cognitive Science* 22(3):269–94. [AJW]
- Wells, G. L. & Gavinski, I. (1989) Mental simulation of causality. *Journal of Personality and Social Psychology* 56:161–69. [aLWB]
- Werner, H. & Kaplan, B. (1963) *Symbol formation*. Wiley. [LM]
- Wiemer-Hastings, K. (1998) Abstract noun classification: Using a neural network to match word context and word meaning. *Behavior Research Methods, Instruments and Computers* 30:264–71. [rLWB]
- Wiemer-Hastings, K. & Graesser, A. C. (1998) Contextual representation of abstract nouns: A neural network approach. *Proceedings of the 20th Annual Conference of the Cognitive Science Society*:1036–42. Erlbaum. [rLWB]
- Wilson, S. G., Rinck, M., McNamara, T. P., Bower, G. H. & Morrow, D. G. (1993) Mental models and narrative comprehension: Some qualifications. *Journal of Memory and Language* 32:141–54. [aLWB]
- Wimmer, H. & Perner, J. (1983) Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13:103–28. [RO]
- Winograd, T. & Flores, F. (1987) *Understanding computers and cognition: A new foundation for design*. Addison-Wesley. [aLWB]
- Wittgenstein, L. (1953) *Philosophical investigations*. Macmillan. [arLWB]
- Wolfe, J. (1998) How quickly they forget: A modest alternative to blinks and blindness. *Abstracts of the Psychonomic Society* 3:53. [JH]
- Wu, L. L. (1995) *Perceptual representation in conceptual combination*. Doctoral dissertation, University of Chicago. [aLWB]
- Wu, L. L. & Barsalou, L. W. (1999) Grounding property listing in perceptual simulation (under review). [arLWB]
- Yeh, W. & Barsalou, L. W. (1999a) A review of situation effects in cognition (in preparation). [aLWB]
- (1999b) Situation effects in conceptual processing (in preparation). [aLWB]
- Younger, B. A. & Cohen, L. B. (1985) How infants form categories. In: *The psychology of learning and motivation, vol. 19*, ed. G. H. Bower. Academic Press. [aLWB]
- Zajonc, R. (1980) Feeling and thinking: Preferences need no inferences. *American Psychologist* 39:151–75. [LCC]
- Zatorre, R. J., Halpern, A. R., Perry, D. W., Meyer, E. & Evans, A. C. (1996) Hearing in the mind's ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience* 8:29–46. [aLWB]
- Zeki, S. (1993) *A vision of the brain*. Blackwell Scientific. [aLWB]
- Zwaan, R. A. (1996) Processing narrative time shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:1196–207. [RAZ]
- (in press) Embodied cognition, perceptual symbols, and situation models. *Discourse Processes*. [RAZ]
- Zwaan, R. A. & Radvansky, G. A. (1998) Situation models in language comprehension and memory. *Psychological Bulletin* 123:162–85. [RAZ]