BEHAVIORAL AND BRAIN SCIENCES (1999) **22,** 577–660
*Printed in the United States of America*

# 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.

0140-525X/99 $12.50

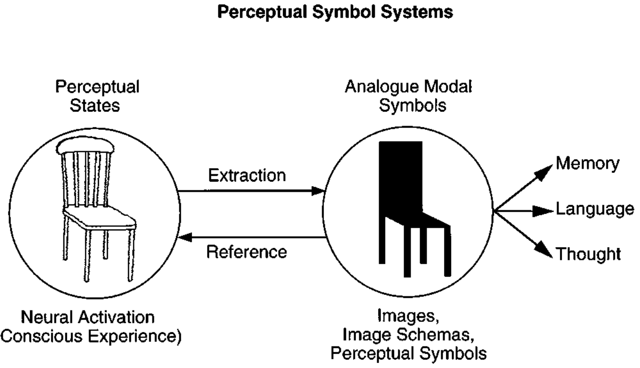

Figure 1. The basic assumption underlying perceptual symbol systems: Subsets of perceptual states in sensory-motor systems are extracted and stored in long-term memory to function as symbols. As a result, the internal structure of these symbols is modal, and they are analogically related to the perceptual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

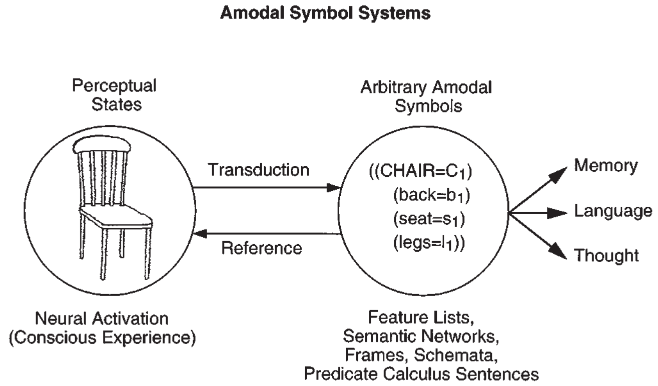

Figure 2. The basic assumption underlying amodal symbol systems: Perceptual states are transduced into a completely new representational system that describes these states amodally. As a result, the internal structure of these symbols is unrelated to the perceptual states that produced them, with conventional 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 & Amhrein 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: "Computational 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 something 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 underlies 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-Molinaro 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 converging 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 representation 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-Molinaro 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-Wilson 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., Baillargeon 1995; Hespos & Rochat 1997). These results further 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 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]

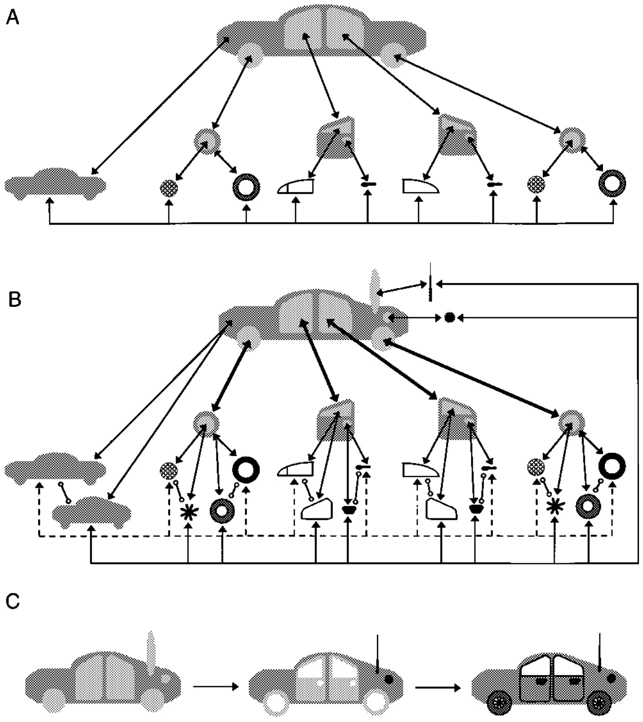

Figure 3. (A) An example of establishing an initial frame for *car* after processing a first instance. (B) The frame's evolution after processing a second instance. (C) Constructing a simulation of the second instance from the frame in panel B. In Panels A and B, lines with arrows represent excitatory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

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, . . . ) 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* (Halff 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 instructions 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 electronics, 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.

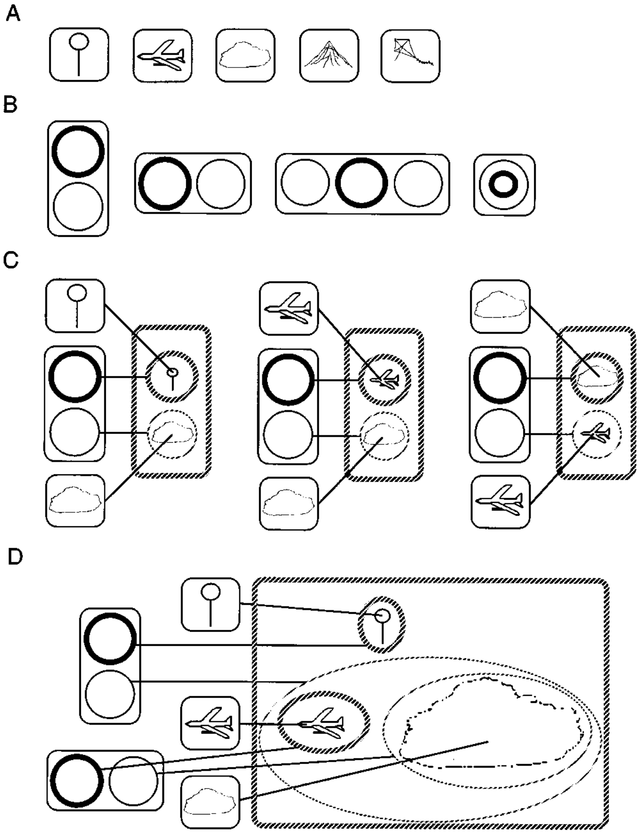

Figure 4. An example of how perceptual symbols for object categories (A) and spatial relations (B) implement productivity through combinatorial (C) and recursive (D) processing. Boxes with thin solid lines represent simulators; boxes with thick dashed lines represent simulations.

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.

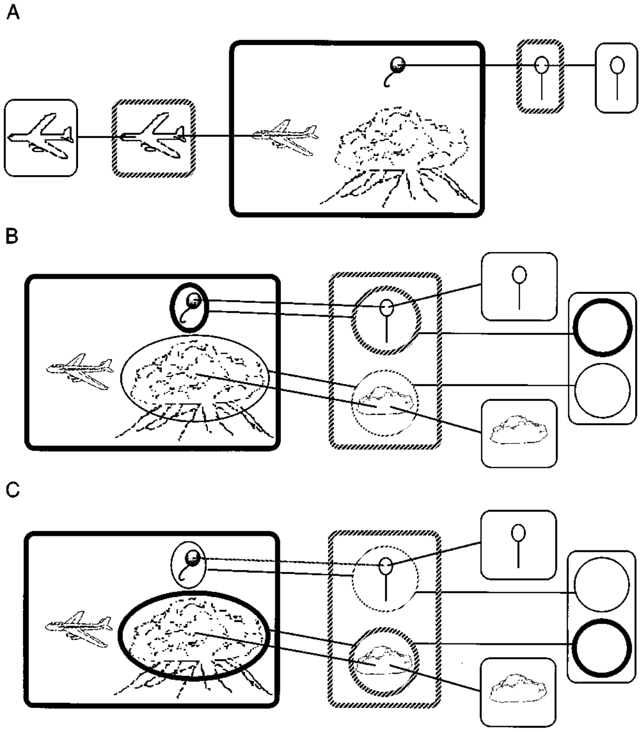

Figure 5. (A) Examples of how perceptual symbol systems represent true propositions by fusing perceived entities with simulations constructed from simulators. (B) Example of a complex hierarchical proposition. (C) Example of an alternative interpretation of the same aspects of the scene in panel B. Boxes 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. Although 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 semicircle. 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 ⊃ 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 tipped. Separating the ⊃ 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-

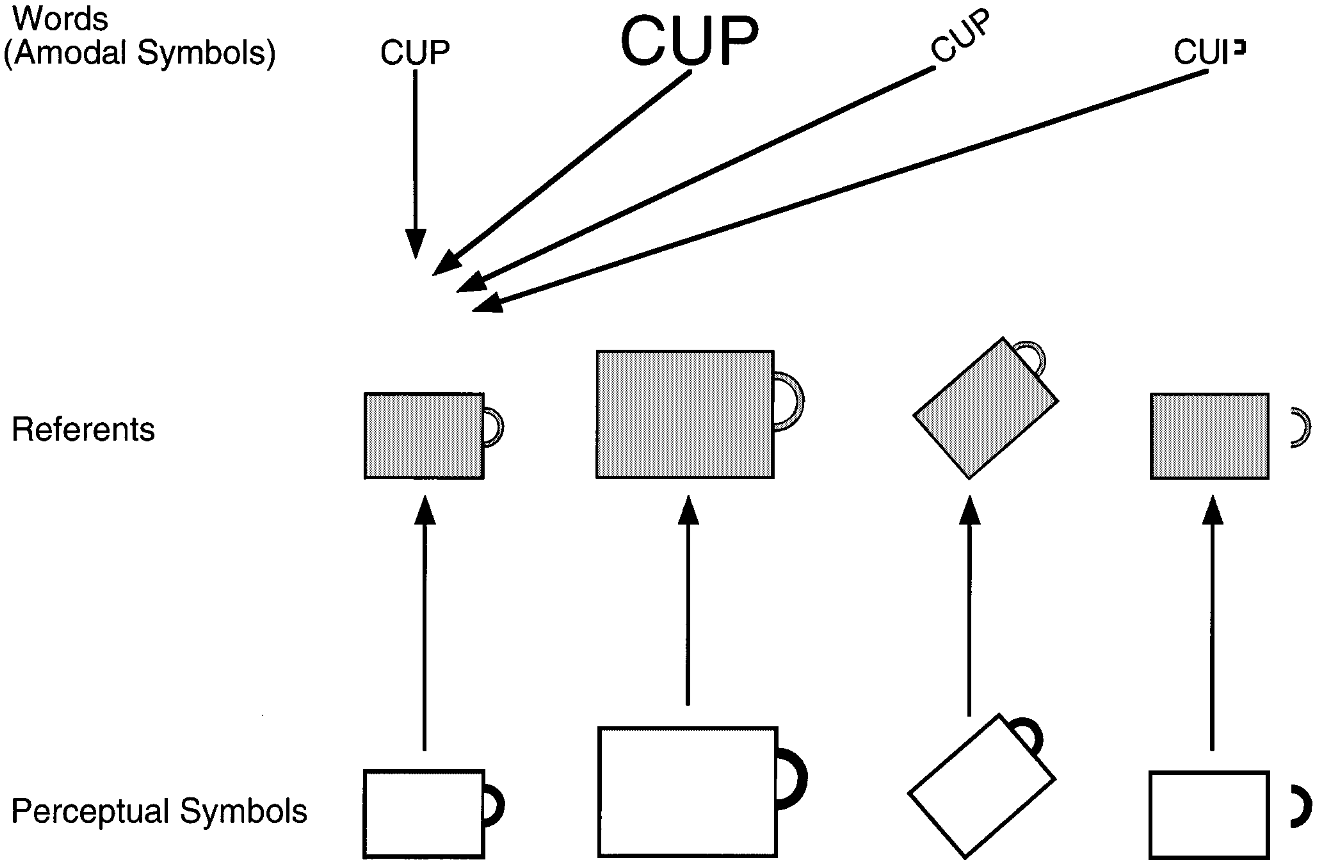

Figure 6. An example of how transforming a word or an amodal symbol fails to produce an analogous transformation in reference, whereas transforming a perceptual simulation does.

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, representing 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, including 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]

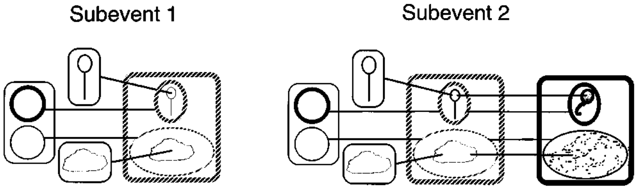

A. "It's <u>true</u> that there is a balloon above a cloud."

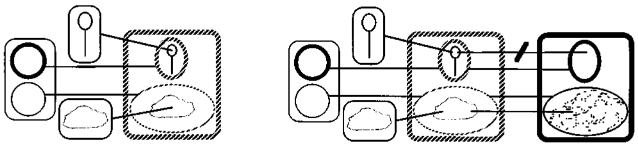

B. "It's <u>false</u> that there is a balloon above a cloud."

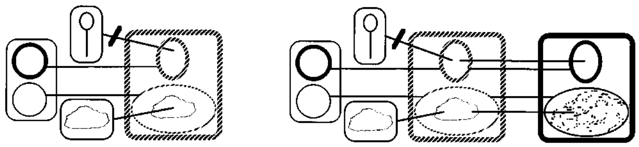

C. "It's <u>true</u> that there is <u>not</u> a balloon above a cloud."

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 simulators; boxes with thick dashed lines represent simulations; boxes with thick solid lines represent perceived situations.

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,* 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.

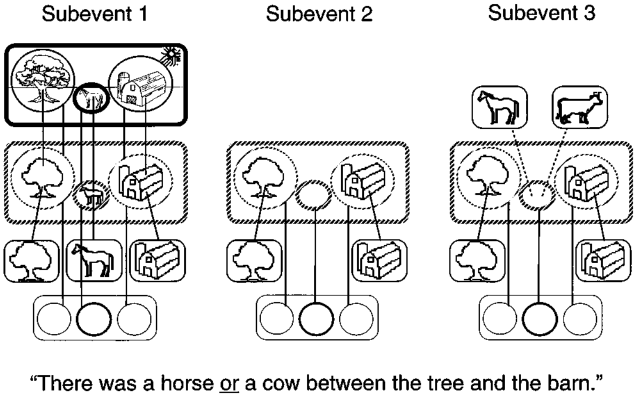

Figure 8. Accounting for one sense of *disjunction* using perceptual symbols. Boxes with thin solid lines represent simulators; boxes with thick dashed lines represent simulations; the box with the thick solid line represents a perceived situation. The dashed lines in Subevent 3 represent alternating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

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

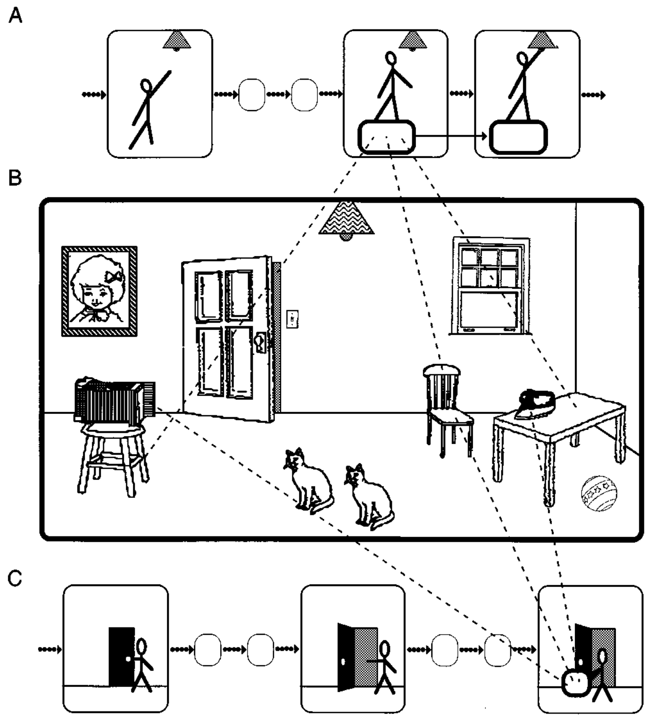

Figure 9. Accounting for the ad hoc categories of *things to stand on to change a light bulb* (A) and *things that could prop a door open* (C), which construe common entities in the same scene differently (B).

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 information 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 reasoning, consider modus ponens. In modus ponens, if the premise X→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→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→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→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→ 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-Molinaro 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, Wenchi 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 theories 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.