Concepts and Meaning

Lawrence W. Barsalou, Wenchi Yeh, Barbara J. Luka, Karen L. Olseth,
Kelly S. Mix, Ling-Ling Wu

University of Chicago

It certainly appears that there should be a relationship between concepts and meaning, but it is not entirely clear what this relation is. We shall assume that *concepts* are people's psychological representations of categories (e.g., *apple, chair*); whereas *meanings* are people's understandings of words and other linguistic expressions (e.g., "apple", "large chair").[1] Currently, many cognitive scientists, especially psychologists, believe that concepts and meanings are at least roughly equivalent, with the meaning of an expression being its conceptual representation in human knowledge. From identifying the content of a concept, the meaning of the associated expression follows. Malt (1991) and Murphy (1991) review this position and various alternatives. In our paper, we shall argue that concepts and meanings differ substantially. Although they are related in important ways, the relationship is one of complementarity, not equivalence.

To reach our conclusion, we shall review standard assumptions about concepts and meaning, challenge these assumptions, and present alternatives. The assumptions that we challenge, and that organize the paper, are: (1) propositional expressions represent concepts; (2) concepts are prototypes of exemplars; (3) concepts are decontextualized and universal in scope; (4) the meanings of words are concepts. We shall argue instead that perceptual symbols represent concepts; concepts are models for types of individuals in world models; concepts are contextualized and local in scope to situations; word meanings use concepts but are not concepts.

Our tack in exploring these issues is to develop a theory of concepts and memory in the first three sections of the paper. In the spirit of cognitive linguistics, this theory utilizes perceptual representations and situational knowledge extensively. The final section compares concepts, as defined in our theory, with relatively well-accepted notions about meaning. We shall then assess whether concepts are equivalent to meaning, or whether they exhibit some other relationship.

Please note that the bulk of this 'working paper' represents a theory of concepts and memory in the early stages of development. We are the first to acknowledge that the majority of our claims lack strong empirical support and that increased precision is necessary at the theoretical level. Although evidence exists for some aspects of our theory, many aspects rest on more of a rationalist analysis of how a cognitive system might compute concepts and meaning. This paper outlines our theory in its current form so that we can begin to examine its claims empirically and implement it computationally. A variety of experimental and simulation projects, both underway and planned, aim to achieve these goals.

Conceptual Representations

The Propositional Approach

Many theories assume that linguistic feature lists represent concepts (for reviews, see Barsalou, 1992b, 1993; Barsalou & Hale, 1993). A feature list contains linguistic descriptions of the characteristics associated with a category's members, such *wings, feathers, beak, flies,* and *builds nests* for *bird.* The simplest interpretation of these features is that they are linguistic expressions in memory.

However, one then needs an account of what constitutes the meaning of the linguistic expression for a feature, simply pushing the problem down a level. The standard move for avoiding this problem is to interpret features as being abstract amodal propositions, represented in some 'language of thought,' such as propositional logic or predicate calculus. So, really, a feature list in memory is not a list of linguistic expressions for features, but is instead a list of descriptions in some form of conceptual representation.

Notably, propositional representations are assumed widely to consist of abstract, amodal, and arbitrary symbols that bear no perceptual resemblance to their referents. On this view, the representation of ABOVE(lamp, table) in the cognitive system consists of symbols that bear no perceptual resemblance to physical instances of *above, lamp,* or *table*. As we shall see, the abstract, amodal, and arbitrary character of the propositional approach leads to significant problems.[2]

Besides feature lists, other, more structured, forms of propositional representation are used widely to represent concepts, including schemata, frames, scripts, mental models, and so forth (Barsalou, 1992b; Gentner, 1989). Much recent work demonstrates the presence of highly structured representations in concepts (Barsalou, 1991; Goldstone, in press; Goldstone & Medin, in press; Markman & Gentner, in press-a,b). Generally, these representations constitute abstract, amodal, and arbitrary 'languages of thought' with origins in predicate calculus. Most recently, connectionist models have provided the newest means of representing concepts, but as we argue elsewhere, connectionist nets are essentially a statistical form of feature lists (Barsalou, 1992b, 1993; Barsalou & Hale, 1993).

Advantages of the propositional approach. Why do cognitive scientists find propositional representations useful? First, propositional representations provide a means for representing the gist of language. Philosophers were the first to point out that paraphrase in language requires that there be some deeper conceptual representation (e.g., Church, 1956), with predicate calculus evolving as the conventional means of representing it. More recently, psychologists realized the importance of propositions after discovering that people have poor memory for the literal or surface forms of utterances (e.g., Anderson & Bower, 1973; Kintsch, 1974; Mandler & Ritchey, 1977; Sachs, 1974). Within 20 seconds or so of hearing or reading an utterance, people typically cannot remember its surface form and only remember its gist. To represent the gist that people remember, psychologists and other cognitive scientists, especially AI researchers, developed a variety of propositional languages, with most being derivatives of predicate calculus (e.g., schemata, frames, scripts, mental models).

A second reason that cognitive psychologists find propositional representations useful is because these languages capture the productive character of language and thought through argument binding and recursion. Argument binding involves mapping the arguments of predicates to individuals, assigning attributes to values, filling slots with instantiations, and so forth, depending on the particular representational scheme. For example, the predicate for *above*, ABOVE(X,Y), contains two arguments, one for an upper region X and one for a lower region Y, which are bound to a variety of individuals across situations, as in ABOVE(lamp, table), ABOVE(bird, tree), and so forth. Recursion results from binding arguments to predicates, enabling infinitely deep structure. For example, the upper region of ABOVE(X,Y) could be instantiated with INSIDE(M,N), as in ABOVE(INSIDE(lamp, box),table). Argumentation and recursion produce the productive quality of many representational systems, from theories of syntax to theories of knowledge. Because an indefinitely large number of individuals can be bound to arguments, and because arguments can often be bound recursively, these representational systems are productive.

A third reason that propositional representations are central to cognitive science is that they are tractable methodologically. We have developed formal logics for manipulating and assessing these representational systems, as well as computer hardware for implementing them. The availability of these tools has in all likelihood defined our theoretical assumptions considerably.

Finally, propositional representations are attractive because natural language provides our most obvious window on human cognition. Although cognitive psychologists often use more indirect measures to draw inferences about cognitive states (e.g., reaction time, percent correct), language provides our most obvious, intuitive, and direct access. Thus, it is perhaps not surprising that our theoretical languages for representing gist are closely related to natural language. Although these representational schemes are really supposed to be languages of thought rather than spoken language per se, they often are remarkably similar to spoken language. As we saw earlier, feature lists are typically expressed in natural language terms. Similarly, predicate calculus expressions in frames, schemata, scripts, and so forth also use language heavily in naming predicates, arguments, and instantiations. Even connectionism typically resorts to naming units (or collections of units in coarse coding) with linguistic names.

Problems with the propositional approach. We believe that serious problems for the propositional approach compromise it significantly. First, we have no accounts of how propositional representations arise in the cognitive system, either through experience or evolution. For example, we have no compelling accounts of mechanisms that transduce propositional representations from perceptual experience. The absence of such accounts is troubling, as is our apparent ignorance of how to establish one.

Second, we have no direct evidence for abstract, amodal, and arbitrary symbols in the cognitive system. Furthermore, we have no direct evidence for the 'left-to-right, shift register' operations that such an account entails (Fodor & Pylyshyn, 1988). Theorists have adopted the propositional approach without obtaining direct evidence for its representational and processing primitives, perhaps because its expressive capabilities serve our theoretical purposes sufficiently well. However, the lack of direct evidence for this approach is again troubling and suggests caution in adopting it.

A third problem is symbol grounding. As a variety of theorists have noted, propositional representations are typically implemented in a purely syntactic manner, leaving their semantics unspecified (e.g., Harnad, 1987, 1990; Searle, 1980). No account is given of how propositional symbols are linked to their referents in the world. Instead, propositions only become linked to the world through the theorists and programmers who create them. When more principled accounts of linking are attempted (e.g., mediating perceptual images), propositional symbols are no longer necessary, because perceptual images can serve their symbolic function, as we describe shortly.

Fourth, propositional approaches exhibit linguistic vagary (Barsalou, 1993). This is the problem of not knowing which aspects of people's concepts to represent in propositional notation. Which predicates and arguments should be included? How deep is recursion? In general, the propositional content that we develop to represent people's concepts can be shown to be unprincipled, haphazard, and incomplete. Content is unprincipled, because we primarily discover it through blind empirical means. Content is haphazard, because it changes with context. Content is incomplete, because further aspects can always be discovered. Our inability to articulate the content of propositional representations in anything close to a compelling manner again suggests caution in adopting them.

In general, the manner in which theorists construct propositional representations seems neither illuminating nor explanatory. Consider the propositional representation of "The lamp is above the table," namely, ABOVE(lamp, table). Essentially, this proposition is the result of dropping the closed class words of the sentence and retaining the content words in a way that captures the conceptual relation between them, using parentheses and a comma. Because a longer string of words has simply been reduced to a shorter string, we haven't actually re-represented the original sentence in a true language of thought--we have simply dropped words that carry information about surface structure and retained words that capture gist. This gist doesn't convey anything conceptual, because it is just words. Again, we have the symbol grounding problem: Reference is not truly established to the world, but is necessarily mediated by the theorist or programmer who constructed the propositional representation.

The Experiential Approach

We believe that the fundamental conceptual representations in the human cognitive system are schematic perceptual images extracted from all modes of experience. The experiential view is in wide currency among cognitive linguists (e.g., Fauconnier, 1985; Lakoff, 1987; Lakoff & Johnson, 1980; Langacker, 1986, 1987; Sweetser, 1990; Talmy, 1988), as well as among other cognitive scientists (e.g., Bickhard, 1980; Huttenlocher, 1973, 1976; Huttenlocher & Higgins, 1978; Mandler, 1992). Indirect psychological support for this view can be found in a variety of empirical studies (e.g., Potter & Faulconer, 1975; Potter, Valian, & Faulconer, 1977; Potter, Kroll, Yachzel, Carpenter, & Sherman, 1986; Glenberg & Langston, 1992; Glenberg, Meyer, & Lindem, 1987; Gernsbacher, Varner, and Faust; 1990). In this section, we present our specific view of perceptual symbols, along with a psychological motivation for this approach in general (see Barsalou, 1993, for a more detailed account).

Basic assumptions. Three critical assumptions are necessary to realize the expressive power that underlies the experiential approach: First, the perceptual representations from experience that represent concepts are not holistic analogue images (i.e., 'pictures in the head'). Instead, they are compositional images in the weak sense of being built analytically from smaller component images.[3] Recent work by Kosslyn and his colleagues supports this assumption, demonstrating that people construct mental images component by component in a systematic manner, rather than retrieving them as single holistic images (Roth & Kosslyn, 1988; Kosslyn, Cave, Provost, & von Gierke, 1988; also see Biederman, 1987).

Second, we assume that perceptual representations are typically much 'sparser' than most actual perceptions. Similar to the construct of *image schema* in cognitive linguistics, we assume that the perceptual representations underlying concepts are abstract and schematic, often excluding much information present in perception. For example, the perceptual representation of *chair* might include components for its seat, back, and legs, while omitting information about color, texture, and so forth.

Third, perceptual representations do not arise solely from vision, nor simply from the five perceptual modalities. Instead, perceptual representations arise from any aspect of experience, including proprioception and the introspection of representational states, information processing operations, and emotions. Consequently, these representations are not 'perceptual' in the traditional sense but are more generally *experiential*, arising from any aspect of experience during perception of the external world and introspection of the internal world. From this point on, *perceptual symbol* refers to a symbol extracted from any aspect of experience, not only perception of the external world.

The construction of perceptual symbols. Consider a general psychological account of how a compositional system of experiential images might develop in the cognitive system. We assume that four stages are central to this process: attention, memory, reference, and composition.

It is well known from substantial work in cognitive psychology that people possess selective attention. As Garner's (1974, 1978) seminal work demonstrated, many aspects of experience are separable. In viewing a triangle, people can focus attention on its shape and ignore its color, or vice versa, such that shape and color are separable aspects of experience. Using attention to select aspects of an experiential image in this manner allows people to represent images componentially.

Once selective attention extracts a component of a perceptual experience, there is good reason to believe that it becomes stored in long-term memory. Following decades of memory research, it is almost axiomatic that focusing selective attention on an aspect of experience establishes a relatively permanent record of it in the cognitive system. If people focus attention on the shape of an object, a record of that shape becomes established in long-term memory. If attention focuses subsequently on the triangle's color, a record of the color is established as a separate memory. Although both memories might be integrated in a common representation (e.g., a blue triangle), the 'whole' is composed of individual memories for the parts. Aspects of the original entity not attended may be excluded, as prior research has shown (Nickerson & Adams, 1979).

Once an aspect of experience becomes stored in memory via selective attention, it can function symbolically, becoming a perceptual symbol. Thus, the shape extracted from a triangle can refer to the entity from which it was extracted, to that shape in general, or to the shapes of other individuals, depending on the semiotic context. Analogous to how a word can refer to entities in the world, so can perceptual symbols.

Finally, perceptual symbols are compositional in the weak sense that they can combine to compose larger wholes. To see this, imagine a triangle. Now imagine a green triangle. Now imagine a pitted green triangle. One way to think about how you represented this series of triangles is that you combined perceptual symbols in a compositional manner. To form a green triangle, you combined perceptual symbols for *triangle* and *green.* To form a pitted green triangle, you further added a perceptual symbol for *pitted.* Note that these perceptual symbols aren't added to a feature list, nor are they values bound to arguments in propositional notation. Instead, the perceptual symbols for *green* and *pitted* are transformations on the perceptual symbol for *triangle*, coloring and pitting it, similar to transformations noted in the imagery literature (Finke, 1989; Shepard & Cooper, 1982).

It is important to add that many concepts may be represented dynamically by a series of perceptual symbols that represent change over time, as in Langacker's representations of verbs (Langacker, 1986, 1987). For example, *bite* can be represented by a series of perceptual symbols that begin with a mouth being closed, opened, and closed on some object, with the entire series being necessary to represent the process. Following Newtson's (1976) work on the breakpoints of perception, we assume that images at certain points in a continuous process carry much more information than images at other points, with only the most informative images being extracted to form a discrete representational series.

The diversity of perceptual symbols. Focusing only on perceptual symbols extracted from vision compromises the viability of the experiential approach significantly. Important expressive power comes from including perceptual symbols from other aspects of experience (e.g., using perceptual symbols from proprioception to represent causality and contingency; Talmy, 1988). Excluding other modalities of experience severely limits our ability to represent concepts from non-visual domains, especially abstract concepts.

The obvious challenge for the experiential approach is to explain abstract concepts. By focusing solely on visual symbols, this becomes difficult, if not impossible, but by including perceptual symbols from other aspects of experience, especially introspection, the representation of abstract concepts becomes increasingly tractable. For example, the determiners *a* and *the* can be represented in terms of the cognitive operations surrounding the retrieval of experiential images from memory. Whereas *a* might mean the retrieval of a generic image from memory, *the* might mean the retrieval of one in particular (e.g., "A dog is a good companion" vs. "The dog followed me on my walk"). Obviously, the semantics of these two determiners is much more complex than these examples illustrate. Nevertheless, it seems possible to describe their semantics in terms of experiential symbols extracted from the retrieval process and the types of images retrieved. Similarly, *truth* can be represented as an equivalence between two represented states of the world, one purported and one perceived. These states can be represented by perceptual symbols extracted from introspective experiences for states of the world, and equivalence can be represented by perceptual symbols extracted from introspective experiences of a comparison process that maps key elements from one state into the other. Finally, many emotion words can be represented by extracting the emotional aspects of introspective states, such as *anxiety* and *anger*.

We make the following conjecture about the representation of abstract concepts: For any abstract concept, it will always be possible to extract a set of perceptual symbols that can represent it. The set may draw from multiple modalities of experience, it may be complex, and it may involve a dynamic process over time, but it will always be possible to find one. For example, the word *kindly* might be represented by a collection of perceptual symbols that represent initial introspective states, followed by the facial expressions, culminating in behavior (cf. Stein & Levine, 1990; Stein, Trabasso, & Liwag, in press). Clearly, the representation of abstract concepts is a major challenge for this view, but the success of cognitive linguists thus far in representing them is encouraging (e.g., Sweetser, 1990; Talmy, 1988).[4]

Perceptual symbols and consciousness. By no means does the experiential view entail that perceptual symbols are conscious. One can envision the extraction of perceptual symbols on computer hardware in the absence of consciousness. Simply assume that physical energy for light, sound, and so forth is digitized by a technological system to produce 'sensory' transductions of the physical world. Perceptual symbols arise in such a system by selecting subsets of these digitized states, storing them, and later using them symbolically and compositionally during intelligent operations. Of course, the actual development of such technology probably lies far in the future, given its departure from standard symbol processing, but at least it is possible to conceive of such systems.

Nevertheless, the initial extraction of a perceptual symbol may be conscious, because consciousness typically accompanies selective attention. Once a perceptual symbol has been used on many occasions, however, consciousness may fall away. As the process of using the symbol becomes increasingly automatic, it may no longer require selective attention for application and function more unconsciously. Consequently, the processing of many perceptual symbols during comprehension and reasoning may proceed outside consciousness. Barsalou (1992a, Ch. 3) provides a brief review of the relations between consciousness, attention, and learning.

Perceptual symbols versus propositions. Earlier we reviewed the advantages and disadvantages of propositional representations. Here we show that perceptual symbols possess the primary advantages of propositional symbols but none of the disadvantages. Perceptual symbols can represent gist. Imagine hearing the sentence, "The lamp is above the table," and representing it perceptually as a schematic image. Later it should be difficult, if not impossible, to determine if the sentence that produced the image was "The lamp is above the table" or "The table is below the lamp." The perceptual representation of the sentence's meaning captures its gist, allowing substantial paraphrase.

Experiential images naturally represent argument binding and recursion, thereby making them productive (Barsalou, 1993). To see this, consider Langacker's (1986) representation of *above* in a schematic image containing two regions of space, one with a higher vertical position than the other. These two regions are arguments in that they can be instantiated by additional images inserted into them, such as schematic images for a lamp and a table. Similarly, recursion can be realized by inserting an image for a relation into a region of another image, for example, inserting an image of IN(bird, cage) into the upper region of an image for ABOVE('upper region', table). All of the structural properties that make predicate calculus a powerful representational language reside naturally in experiential images as well. Not only can the experiential approach account for gist, it can also account for the complex productive structure of human knowledge.

While maintaining these benefits of propositional representations, experiential images do not exhibit the disadvantages. Although we have no account of how propositional representations originate, and although we have no direct evidence for their predicate calculus-like structure, we can provide plausible accounts for the origin and structure of perceptual symbols: These symbols are literally extracted from perception, and their structure reflects the structure of perception, as well as additional structure imposed through selective attention and memory. Although these stories need to be fleshed out considerably, it is obvious that plausible stories can be told, unlike the case with propositional representations.

Also unlike propositional representations, experiential representations do not exhibit the problem of symbol grounding. Perceptual symbols can obviously be grounded in the entities from which they were extracted, as well as from perceived entities that resemble them.

Finally, perceptual symbols explain linguistic vagary, namely, the problem that linguistic descriptions of concepts are unprincipled, haphazard, and incomplete. Linguistic vagary simply reflects the maxim that a picture is worth a thousand words, or in more technical terms, an experiential image can be described by an infinite number of linguistic descriptions. For each of the infinite possible aspects of an image to which selective attention could be applied, there is a potential linguistic expression that describes it. For this reason, the linguistic descriptions of concepts are unprincipled, haphazard, and incomplete.

Feature lists revisited. We began with a consideration of feature list representations of concepts. Our brief excursion through perceptual symbols suggests a somewhat different way of thinking about feature lists. Rather than being a list of propositions stored in memory, a feature list reflects the sequential description of an experiential image. When people define a concept, they retrieve or construct a schematic image, focus attention on a subset of its perceptual symbols in a sequential manner, and describe the content of each focus with a linguistic description. On this view, feature lists do not exist in long-term memory as conceptual representations but are the result of a sequential on-line process that describes experiential images.

Earlier we noted that one advantage of propositional representations is their potential for being implemented on current computer hardware. We also noted that language provides perhaps the most obvious window on the cognitive system. If the experiential view is correct, then we have to reevaluate these two observations. First, we need to realize that our current computer hardware doesn't come close to providing the kind of representational medium needed to represent human concepts. Instead, we need to develop a completely new type of 'perceptual computer' and 'perceptual syntax,' if our goal is to build a machine that processes information like a human. Second, we need to realize that language does not provide a direct window on human concepts. Instead, concepts exist in a form considerably different from language and its propositional cousins. Concepts consist of perceptual symbols, which language, through the sequential operation of selective attention, describes in a relatively unprincipled, haphazard, and incomplete manner.

In contrast to language, gesture may provide a much better window on concepts. Following the work of McNeill (1992) and Goldin-Meadow, Alibali, and Church (1993), there are reasons to believe that gesture expresses certain aspects of perceptual symbols directly, unlike spoken language. Because gesture is capable of expressing perceptual symbols analogically, it may provide an extremely useful methodology for studying them.

Conceptual Ontology

We turn to the ontological assumptions that underlie theories of concepts, namely, assumptions about the relevant types of entities. To establish our ontology, we develop the constructs of *individual*, *model*, and *world model*, all of which build on the construct of *perceptual symbol* developed earlier. We also develop the construct of *frame*, together with the *one-entity one-frame principle*.

The Statistical Metaphor: Exemplars and Prototypes

According to one common framework, the relevant ontological entities are exemplars and prototypes. An exemplar is typically viewed as the memory trace established in the cognitive system after encountering a member of a category. On encountering a particular dog, for example, a trace of this experience becomes established in memory, thereby establishing an exemplar of *dog*.[5] A prototype, on the other hand, is typically viewed as the central tendency of a category's exemplars, where the central tendency is most frequently computed as a mode or weighted average. For example, the prototype for the category *dog* might be the modal properties across experiences with particular dogs. On reflection, it is clear that this ontology of concepts rests on the statistical metaphor of cases and central tendency. Similar to descriptive statistics, exemplars constitute cases, and prototypes constitute central tendencies across them.

The World Model Metaphor: Individuals and Models

The metaphor of a world model provides an alternative ontology for a theory of concepts (Barsalou, 1991). A world model is a person's beliefs about the current state of the world. It is not beliefs about the types of things in the world, such as taxonomic knowledge about birds and tools; instead, it is beliefs about particular individuals in the world, along with their current states and locations.

We propose that the core of a world model is a hierarchical system of spatial frames, much like the sort of system envisioned by Minsky (1977). This spatial core represents locations in the world hierarchically, from large regions, such as continents and oceans, to smaller regions embedded within them recursively, such as countries, states, cities, neighborhoods, houses, rooms, drawers, and so forth. Each person's knowledge of locations is at least somewhat idiosyncratic, but in each case, a person represents a collection of locations in the world hierarchically. Much remains to be learned about these hierarchical systems of spatial locations, but it wouldn't be surprising to discover that they are represented perceptually rather than propositionally.

Within the spatial core of a world model, we propose that people represent the current positions and states of familiar individuals. Within a home, someone might represent individual chairs, plants, and pets. Similarly, in an office, someone might represent a computer, desk, and books. In representing an individual, people not only represent its location but also its current state. In an office, someone might represent the computer as currently running. Besides representing individuals, world models also represent events. While at work, one might believe that a committee meeting is taking place down the hall, that a plumber is clearing a drain at home, and that Congress is debating a bill in Washington.

Besides representing the current state of the world, world models may also organize information about the past and future, providing the general organization of memory. For example, the representation of a home in someone's world model might specify the color of its walls at some point in the past prior to a recent painting. Similarly, it might represent a prior arrangement of furniture, or a light fixture that once existed but has now been removed. On this view, a world model stores previous states of the world within the same spatial system as current states. Similarly, envisioned future states of the world may be represented as well. If a person is planning to paint the walls of a room or to rearrange its furniture, these future states are stored in memory at their respective locations in the world model. Events from the past, present, and future that occurred in the same location are stored together, with the most recently and frequently processed events dominating.

A variety of temporal structures may organize the information in world models orthogonally to the hierarchical spatial system at its core. Although much remains to be learned about the nature of these structures and their roles in cognition, it seems reasonable to assume that such structures include the daily, weekly, monthly, seasonal, and yearly cycles, as well as more content-oriented time lines for daily routines, life span development, formal education, and so forth (Barsalou, 1988; Conway, 1990). Although we have little to say about these structures or their use, we assume that they organize past, current, and future information in world models orthogonally to space.

Individuals. In the scheme of a world model, individuals, both animate and inanimate, are central ontological entities. Individuals might seem analogous to exemplars in the statistical scheme, but they differ in important ways. Individuals are not memory traces but are continuously existing entities in the world. Because the same individual may be processed on many occasions, many different traces may be produced for it by the cognitive system. Rather than existing independently of one another, as in an exemplar model, these cases are merged together to represent the continuous individual over time. Consequently, the representations of individuals are not purely episodic, as are traces in an exemplar theory; instead, they contain both generic and episodic information.

More specifically, we propose that *frames* represent the individuals in a world model. A frame is a data structure that serves to capture stable information about the individual over time, as well as variability (Barsalou, 1992b; Barsalou & Hale, 1993). As developed in Barsalou (1993), frames are large collections of perceptual symbols integrated to form a unified representation of an entity, in this case, an individual. For example, the frame that represents a particular car, say one's own, begins with a collection of perceptual symbols for its known aspects, such as wheels, color, body, engine, and radio. These perceptual symbols are organized spatially to represent their configuration in the individual. As new aspects of the car are discovered, they are added to this stable configuration (e.g., floor mats). Over time, variability in these aspects may be noted. For example, the gas tank may be full or empty, the body surface may be clean or dirty, the antenna may be present or absent. In each case, the alternative states constitute specializations of an aspect in the perceptual frame; or in the more standard language of frames, these specializations constitute different values of the same attribute. For example, *clean* and *dirty* constitute two specializations of the aspect *surface condition*. The association of specializations to the aspect they specialize organizes this variability into a conceptual field of contrasts (Barsalou, 1992b). Constraints between different types of specialization further organize variability, such as the correlation between tire pressure and steering precision.

On this view, the frame for an individual is neither a single memory trace nor a purely episodic representation. It is not a trace, because many memory traces are organized into it. Across encounters with one's car, traces for different aspects noted on different occasions become assembled into a common frame for the individual. However, the frame for the individual is not purely episodic, because an individual aspect, such as the steering wheel, may be noted on many occasions such that the repeated processings produce a generic representation whose episodic origins become lost.

Furthermore, the representation of an individual may be viewed meta-cognitively as generic to a large extent. Because the goal of the cognitive system is to build a general representation of what the individual is like, the frame for the individual is viewed, not as representing the individual on a particular occasion, but as representing it in general. Across new encounters with the individual, constant aspects in its representation are verified and strengthened, while new aspects are added, either as aspects never noted before, or as specializations of existing ones.

Although much of an individual frame is generic, it always maintains an episodic component, namely, the specializations extracted from the respective individual on the most recent encounter with it. For example, the most recent reading of the gas gauge in one's car is an episodic specialization of that frame aspect, as is the current dirtiness of the surface and the last setting of the radio dial.

The one-entity one-frame principle. As just illustrated by our discussion of individuals, we believe that one and only one frame represents each individual. Even though an individual may be encountered on many occasions, with diverse information extracted on each, all of this information is integrated into a single frame that represents the individual. Thus, the representation of individuals constitutes a special case of the one-entity one-frame principle, where the entities are individuals.

By following the one-entity one-frame principle, frames for individuals represent two important ontological assumptions implicitly: First, each frame implies that there is a single corresponding individual in the world. Second, each frame implies that its respective individual exists continuously over time, because information from multiple events is integrated into a continually evolving structure. Certainly, individuals have temporal beginnings and endings, but within these bounds, a frame represents the continuous existence of a single individual.

Shortly, we shall extend the one-entity one-frame principle to models, proposing that each frame for a model represents a single *type* of individual. This particular extension will be central to our later definition of concepts, and to our later comparison of concepts and meaning.

Models. We propose that people construct models for significant types of individuals. We further propose that people view models as having a different ontological status than individuals. Whereas an individual is a physical entity believed to exist continuously in the environment, a model is a mental construct for a type of individual. Unlike individuals, models do not have corresponding physical entities in the world. For example, people do not believe that a physical model of *dog* exists in the world that corresponds to their internal model of *dog*. Certainly, different individuals in the world may instantiate a model, but no direct physical counterpart to the model typically exists. Instead, people view their models for types as only existing mentally.

It is important to distinguish our use of *model* from the related constructs of *world model*, *prototype*, *mental model*, and *situation model*. Models differ from world models in that a model represents a type of individual, whereas a world model represents the current state of the world. Models differ from prototypes primarily in how they are updated (cf. Smith & Medin, 1981): First, models, as we shall see, are updated locally, only being revised with respect to the most recently encountered individual, whereas prototypes are revised globally with respect to all previous exemplars (for a similar view, see Medin & Ross, 1989; Ross, Perkins, & Tenpenny, 1990; Ross & Spalding, 1991). Second, we shall propose that models are often revised hierarchically, through the specialization of perceptual symbols, whereas prototypes are revised by computing averages or modes across flat lists of features. Finally, models differ from mental models and situation models, first, in representing generic types, not episodic situations, and second, in being perceptual representations, not propositional (cf. Glenberg et al., 1987; Johnson-Laird, 1983; van Dijk & Kintsch, 1983).

The literature on generics documents linguistic devices that mark individuals and models, consistent with our claim that people distinguish them ontologically (e.g., Carlson, 1982; Schubert & Pelletier, 1987). Consider some examples: In "I don't understand my computer," "my computer" refers to inadequate understanding of an individual computer. Similarly, in "The computers in the front office are slow," "the computers" refer to multiple individuals. However, in "I have no idea what a computer does," "a computer" refers to an inadequate model of computers in general. Similarly, in "This machine is too much computer for my needs," "too much computer" again refers to the model of a computer. The fact that languages mark individuals and models explicitly in this manner suggests strongly that human knowledge represents both and distinguishes them ontologically.

We assume that models, like individuals, are represented by frames. Furthermore, we assume that frames for models constitute another special case of the one-entity one-frame principle, where the entities are models. Specifically, we propose that people develop one and only one frame for each important type of individual that they believe exists in the world. Thus, each frame for a model implicitly represents the ontological belief that a corresponding type of individuals exists, has existed, or could exist in the world.

Analogous to frames for individuals, the frame for a model is a data structure that integrates perceptual symbols across multiple individuals, capturing constancies and organizing variability. Also like frames for individuals, frames for models contain both generic and episodic information, with generic information comprising repeated perceptual symbols across individuals, and with episodic information being their most recent specializations from last-processed individuals. The next two sections provide greater detail about frame creation and frame revision for both individuals and models.

Frame creation. A critical issue arises in our framework: On encountering an entity in the environment, should the information extracted from it be assimilated to an existing frame in memory, or should a new frame be created to accommodate it? We propose that the one-entity one-frame principle guides this process. To see how, first consider this issue at the level of individuals. If an entity is not recognized as a familiar individual, then a new frame for an individual may need to be constructed, because each new entity requires a new frame. However, we frequently encounter new entities for which we are unlikely to construct individual frames. On walking down a busy street, we probably don't create a new frame for every stranger that walks by. Instead, these individuals only exist perceptually, and possibly in working memory, with frames for them never becoming established in long-term memory. Similarly, on driving down the freeway, we probably don't create a new frame for every car that passes us, or for every tree that we see. Because these are not significant individuals that are likely to be central in our lives over an extended time period, there is no need to establish a permanent frame for them in long-term memory. Instead, they can be represented in perception and can be understood using the model frame for that type of entity. We propose that people only create a new frame for an individual when the individual is both unrecognized and expected to be significant. For example, buying a new car should establish a new frame, because this individual is likely to be significant for some time. Similarly, meeting a new neighbor or planting a tree in the front yard should lead to new frames for these individuals.

A variety of cues may suggest the presence of a new individual. First, failing to recognize the individual as a known individual may suggest a new one, such as failing to recognize a person. Second, another person may describe an entity in a way that introduces it as a new individual, as when a friend says, "Let me tell you about my new car." Third, the presence of an entity in a particular location may signal that it is novel, such as a tree in a spot that was previously empty. In general, there may be a wide variety of cues that allow people to recognize familiar and unfamiliar individuals.

Once frames are established for individuals, they should typically be accessed when the individuals are encountered on subsequent occasions. Following the one-entity one-frame principle, new frames should not be created. Certainly, a known individual may not be recognized under certain circumstances, as when the individual is still relatively novel, when it undergoes a major transformation, or when there are highly similar individuals in the vicinity. Assuming, however, that the individual is recognized correctly, information extracted from it on this and subsequent occasions is integrated into its already existing frame, following the revision procedure described shortly.

Now consider the creation of frames for models. Again, many cues may be relevant to determining that an entity is an instance of some new type of thing never encountered before. In these cases, the one-entity one-frame principle prescribes that a new frame for this type be created. For example, an animal may exhibit a configuration of parts that is sufficiently novel to prevent its being assimilated to any existing model of an animal. Similarly, someone might introduce an entity as being of a novel type, as in saying, "Here's a tool I'll bet you've never seen." Again, the new model must be of sufficient importance for a new frame to be established. For example, many people can't tell the differences between birches, beeches, and other similar trees, presumably because it's never been important to establish separate models for them. Instead, these types of tree are all assimilated to the more general model for *tree*, or perhaps to the model of a specific type of tree for which a model exists, such as *oak*.

Although many models may represent types of things, such as types of plants, animals, substances, and artifacts, models may often be constructed for other reasons. For example, people may construct a model to represent a set of entities relevant to a particular goal, such as *foods to eat on a diet* (Barsalou, 1991). Although this category is not a taxonomic kind, people may nevertheless want to accumulate generalizations about it in a model that could be useful to the goal it serves. For example, integrating information about calories, food group, and taste across instances may be useful for recognizing and using *foods to eat on a diet* .

Frame revision. When a familiar entity is encountered, new frames should not be created, according to the one-entity one-frame principle. Instead, existing frames should be revised. If an individual frame exists for the entity, it should be updated to reflect the entity's current properties. Similarly, if a model frame exists for this type of entity, it should be updated, too. Although we only discuss the revision of model frames, we believe that the revision procedures presented in this section operate on the frames for both individuals and models.

When the frame for a model is first constructed, it could contain all of the perceptual symbols extracted from the initial individual; alternatively, it might only contain a subset, depending on the amount of time available to process the individual, whether certain aspects of the individual are believed to be idiosyncratic, and so forth. On encountering a second individual later, two things may happen: First, if some aspect of the second individual is not represented in the model, a perceptual symbol for this aspect may be added, assuming that it is not believed to be idiosyncratic. Second, if some aspect of the second individual contrasts with an existing aspect in the model (acquired from the first individual), the *superordinate aspect* of the two aspects is computed and stored in the model. Additionally, the new aspect from the second individual specializes the superordinate aspect, and the aspect from the first individual remains associated to the superordinate aspect as a contrasting specialization.

To see this, imagine encountering an entity that is a new type of kitchen tool. Because you believe that this is a significiant type of entity, you establish a new frame to represent its model. Initially, the model frame only contains perceptual symbols extracted from this first individual, such as its handle. On encountering a second instance of this tool that exhibits a differently shaped handle, a superordinate aspect that captures both handles may become established in the model. The superordinate aspect is a perceptual symbol that loses information in order to accommodate both specializations. For example, if one handle is straight and thin, whereas the other is curved and thick, their superordinate aspect may be a general cylindrical shape whose thickness is intermediate.

Once the perceptual symbol for a superordinate aspect has been constructed, the most recent specialization is inserted into it, and the previous specialization remains associated as a contrast. In the handle example, the superordinate aspect is specialized to be curved and thick, and the previously extracted, straight and thin handle remains associated. Although the model becomes more general by adding a superordinate aspect, it remains current by maintaining the most recent specialization, and it organizes variability by contrasting related specializations.

Certainly, much remains to be learned about the process of constructing perceptual symbols to represent superordinate aspects. One possibility is that superordinate aspects tend to minimize complexity and seek optimal goodness of form. For example, if *straight* is less complex and has better goodness of form than *curved*, the superordinate aspect that accommodates them may evolve toward being more straight than curved (e.g., Tversky, 1991). Alternatively, superordinate aspects may become increasingly 'ghostly' as they evolve to cover increasing numbers of specializations. As Huttenlocher (1973) notes, aspects of images often seem quite amorphous, omitting color, texture, and so forth, as they lose information originally in perception.

In many cases, a superordinate aspect may primarily serve as a place holder, representing an *instantiated region*, rather than capturing essential properties common to all specializations. In the handle example, the superordinate aspect that represents multiple handles may not attempt to capture common features, nor be a perfect average of its specializations. Instead, the superordinate aspect may primarily serve to indicate the region where some kind of handle exists and varies. A more precise notion of what a handle is can be obtained, if necessary, by consulting the region's specializations.

Transcendence. Thus far, we have assumed that the frames for individuals reside at specific locations in a world model. However, some individuals change location often, such as one's spouse or car. At a given point in time, an individual is only at one location. Over time, however, the individual may move from location to location. How does the cognitive system represent such individuals over time? A similar problem concerns the representation of models: Because models do not correspond directly to any individual in the world, where do they reside in a world model, if anywhere?

We propose that the frames for individuals and models become increasingly transcendent of world model position to the extent their positions change. Whereas frames for individuals and models that occur regularly at a particular location become grounded in that location, frames for individuals and models encountered in many locations establish high degrees of transcendence. We are not proposing that transcendence constitutes a separate memory store from the world model. Instead, we propose that transcendent frames become detached from the world model *functionally*, while remaining associated structurally to many locations.

When a frame is created initially, either for an individual or model, it becomes associated with the location in the world model currently salient. If an individual frame for a new car is created on first seeing it at a dealership, the frame becomes stored at the dealership in the world model. Similarly, if the car were the first instance of a new type (e.g., a newly introduced line of cars), the frame established for the model would be associated initially with this particular dealership as well. As these frames are imagined in additional locations, they become increasingly transcendent. If the individual car were further imagined in one's garage (e.g., to simulate how it would fit there), its frame would become associated with two locations: the dealership and garage. Similarly, if one started seeing additional instances of this new line of cars around town, the frame for the model might become associated with these additional locations in the world model.

The increasing transcendence of a frame causes it to be retrieved in isolation, rather than in a particular location. Although the frame remains structurally associated with all of its locations, these multiple associations produce interference that 'cancel each other out,' thereby producing functional transcendence at retrieval. Evidence for this canceling-out phenomenon comes from Thorndyke and Hayes-Roth (1979) and Watkins and Kerkar (1985), who demonstrated that idiosyncratic information in related events becomes increasingly inaccessible, whereas shared information becomes increasingly accessible. Extending this phenomenon to frames in world models, common information in a frame across locations becomes strengthened, whereas idiosyncratic information about locations cancels itself out.

When a frame is associated with many locations, and when the associative strength of no one location is sufficiently strong to overcome the competitive interference from all other locations, the frame is retrieved independently of a location and transcends the world model. For example, the model frame for *car* may be associated with so many locations, none of which dominate, that it becomes transcendent. Conversely, a location for one's own car, such as a home garage, may be sufficiently dominant that the frame for this individual is often retrieved in this location. Although this individual frame remains associated with many other locations, the dominant location is sufficiently strong to preclude transcendence. Should the retrieval context favor a less dominant location, the frame could be retrieved in that context instead, such as thinking about the last time the car was serviced and retrieving the individual frame in the location of an automotive shop.

Frames for models may similarly become transcendent to the extent that the individuals producing them occur across many locations. As one encounters many instances of a new line of car across many locations, its model frame may become increasingly transcendent. Again, however, if one location dominates, the model may typically be retrieved in it. For example, if one only see instances of a particular type of car at a particular dealership, its model frame may typically be retrieved in this location.

As individual and model frames become transcendent, they probably do not 'float around' in memory disorganized. Instead, they may become organized into transcendent taxonomies and other sorts of decontextualized conceptual fields. These transcendent organizations may serve at least the following two functions: First, they may organize general beliefs about the kinds of things in the world, capturing important similarities between models, as well as providing useful inferences, such as inferences from cognitive economy (e.g., inferring that zebras have hearts from knowing that mammals have hearts). Second, these transcendent organizations may provide the building blocks of world models (Barsalou, 1991). When a new individual is encountered, an appropriate model is retrieved from a taxonomy, and a copy of the model is inserted into the world model to represent the individual (analogous to how copies of generic objects are copied into drawings from the palette in object-oriented drawing programs). On encountering an unfamiliar dog, for example, the model for *dog*, or perhaps a more specific breed, is retrieved, copied, and inserted into the world model to represent the individual.

This copy-and-insert process serves two functions: First, it is not necessary to process an unfamiliar individual extensively to represent it. As soon as the individual is recognized as being an instance of a model, a copy of the model can be used to represent the individual in the world model. This process greatly reduces the amount of learning about the individual via bottom-up processing that must be done to represent it. In the limit, the current form of the model may simply represent the individual without the addition of any perceptual information from the individual, should there be insufficient time or motivation for specializing it. The second purpose served by copying the model to represent the individual is that the copy may carry information not available in perception. For example, the engine of a perceived car may not be perceptible, but copying the model for *car* to represent it could easily provide default information about the likely presence of one.

Conceptual Situations

Theories of knowledge often assume that concepts are context-independent and universal. Concepts are context-independent when they represent exemplars in isolation, omitting the typical situations in which they occur. For example, a context-independent concept for *chair* might only represent the physical parts of chairs, omitting the situations in which they are normally found, such as a library or living room. Concepts are universal when they attempt to cover all relevant exemplars simultaneously. For example, a universal concept for *chair* might attempt to provide a set of features that identifies every possible chair in the world and excludes all non-chairs.

To the contrary, we shall argue that people's concepts are neither context-independent nor universal but are situated and local. By situated, we mean that people typically represent concepts in the larger situations that contain them. For example, we shall argue that, typically, a person's concept for *chair* represents chairs in particular situations, such as chairs in a living room. By local, we mean that the concept constructed by a person on a particular situation is local to that situation, only covering exemplars in that situation and not attempting to cover all exemplars universally. For example, we shall argue that, typically, a person's concept for *chair* only covers those chairs in the current situation containing it, such as the chairs in a particular living room.

In the previous section, we introduced the constructs of *individual* and *model*, characterizing both as involving frame creation initially and frame revision subsequently. Actually, the representation of individuals and models is much more complex than simply creating and revising a single frame for each. Once situations are introduced, *specialized frames* develop within each individual and model frame, for each of the typical situations in which the frame occurs.

In this section, we extend the perceptual frames developed in the previous two sections to the representation of situations. We define what we mean by *situation* and then discuss two basic types: episodic and generic. Whereas an episodic situation represents a single event that occurred at a specific time, a generic situation generalizes over related episodic situations. The distinction between episodic and generic situations is orthogonal to the distinction between individuals and models: The representation of an episodic situation may contain both individuals and models, as may the representation of a generic situation.

Situations

Our construct of *situation* is similar in spirit to Johnson-Laird's (1983) notion of a mental model, although we propose explicitly that situations are composed exclusively of perceptual symbols (see Glenberg et al., 1987, for a similar view). Our construct of *situation* is also similar, although not identical to, Lakoff's (1987) notion of an idealized cognitive model. One difference is that our situations are not necessarily idealized but instead represent what occurs normally in situations (Barsalou, 1985). Finally, our emphasis on situations is related to recent work on situated cognition (*Cognitive Science*, 1993, Volume 17, Number 1). In contrast to this work, which is fairly anti-cognitive, we believe that providing computational accounts of situations is essential.

In our framework, events are composed of situations, which in turn are composed of images. Note that *event* and *situation* in our terminology refer to cognitive representations, as does *image*, not to the physical world. Events, situations, and images parallel, at least somewhat, the constructs of scripts, scenes, and states in Schank (1975, 1982) and Schank and Abelson (1977). Because we represent these constructs with images and perceptual symbols, whereas Schank and Abelson represent them with propositions, our accounts differ considerably in many ways associated with these alternative forms of representation.

We define an *image* as:

(1) a set of perceptual symbols,
(2) representing individuals and/or models,
(3) in a static spatial configuration,
(4) perceived from a particular perspective.

For example, an image might be a frontal view of flowers in a vase on a table against a wall in a room.

We define a *situation* as:

(1) a series of images,
(2) depicting a relatively constant set of individuals and/or models,
(3) changing in some significant way continuously over time,

(4) in a relatively constant region of space.

For example, a situation might contain a series of images in which a person puts a vase of flowers on a table. An initial image might depict a person next to a table holding a vase, followed by images of the person placing the vase on the table, stepping away from the table, leaving the vase on top. As this example illustrates, the individuals remain constant in a constant region of space, with their configuration changing to represent a significant event, which might be the changed location of the vase, or the presence of something new on the table.

An important issue concerns the representation of continuous change over time in a situation. In principle, an infinite number of images are necessary to represent a situation continuously. We suspect, however, that people store only the most informative images within a situation, those receiving their greatest attention. As demonstrated by Newtson (1976), people reliably perceive salient 'break points' in perceived event sequences, where these points can generally be construed as occurring after major qualitative changes in the configuration of individuals. Because the cognitive system may often be able to simulate the likely path between the images at two adjacent break points, it may only store images at break points when representing the situation, computing intermediate points on line when necessary. For example, the images of a person next to a table holding a vase, placing the vase on the table, and then stepping away may capture the break points of this situation, because they represent qualitatively different configurations of individuals, and because the intermediate images can be simulated easily.

Finally, we define an *event* as:

(1) a series of two or more situations,
(2) related in a coherent manner,
(3) leading to a significant outcome.

For example, an event might contain several situations that culminate in a vase of flowers being placed on a table. A first situation might depict cutting flowers in a garden, a second walking from the garden to the kitchen, a third putting flowers in a vase, a fourth walking to the living room, and a fifth placing the vase on the table. As this example illustrates, the individuals change, at least somewhat, across situations, as do the regions. However, the situations are coherently related, culminating in a significant outcome (cf. Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985; Trabasso, van den Broek, & Suh, 1989).

Episodic Situations

Like images, situations, and events, episodic situations are cognitive representations of the world, not physical events. Prototypically, an episodic situation results from perceiving a physical event in the world, with the episodic situation being a representation of a single situation within the physical event. Unless the perceiver is lost in the world, the episodic situation is stored in his or her world model at the location in which the physical situation occurred. Additionally, the episodic situation is associated with temporal structures, such as those representing time of day, daily routines, life periods, and so forth.

Less typical episodic situations also occur. In these cases, episodic situations aren't experienced perceptually but are imagined. In counterfactuals, a person constructs an episodic situation that didn't occur but is an alternative to an episodic situation that did. Typically, these alternatives share the same location in the world model and the same temporal indices but vary in the individuals they contain, the states of these individuals, or the relations between them. Similarly, in planning future events, a person constructs an episodic situation that hasn't occurred yet but that could. Often, these 'prospective memories' are associated with specific locations in a world model and specific indices within temporal structures, but not necessarily. We exclude planning for a basic *type* of situation, given that these cover multiple episodic situations and are therefore generic situations.

Mixture of episodic and generic information. Although imagined episodic situations are interesting and important, we focus primarily on experienced episodic situations. Contrary to an implicit assumption in most exemplar models, we do not believe that experienced episodic situations are entirely episodic--they are not 'pure' memory traces of physical events. Instead, we propose that episodic situations contain substantial amounts of generic information as well. As a physical event is experienced, frames for individuals and models relevant to the event become active and are associated with the current location and time in the world model. New frames may also be created if they are perceived as necessary for representing new, significant individuals or new types of things in the physical event. Perceptual information extracted from the physical event then specializes these individual and model frames (when perceptual symbols already exist that can be specialized), or this extracted information is added to these frames (when new perceptual symbols must be added to represent it). The current forms of all active frames are then associated with each other, as well as with location and time in the world model. Consequently, an episodic situation is a mix of both episodic and generic information: The generic information includes pre-existing knowledge of individuals, models, location, and time, whereas the episodic information includes specializations of frame information and additions to it, along with associations that integrate individuals, models, location, and time.

To see this, imagine watching one's spouse place a new vase containing flowers on a table in the living room just before dinner.[6] During the physical event, frames for known individuals become active, including one's spouse and the table, as well as model frames for the basic level categories to which they belong, *woman* and *table.* Note that both individual and model frames become active for familiar individuals, an assumption that will be relevant later to the construction of generic situations. Models for unfamiliar individuals also become active, such as the models for *vase* and *flowers*. Finally, frames are created for any new individuals and models that are significant. Because the vase is new and potentially significant (i.e., it is likely to become a fixture), the model frame for *vase* is copied into this location of the world model to create a new individual frame for this particular vase, with the copy possibly being specialized by information extracted from perception. Although the flowers are also unfamiliar, a new individual frame may not be constructed for them, assuming that they are not significant (i.e., they are likely to be gone in a couple of days). Instead, a specialization of the model for *flowers,* or perhaps for the specific type of flower (e.g., *tulips*), represents them in the episodic event. If the flowers were of a type never encountered before, and if the perceiver takes an interest in flowers, he or she might construct a new model frame for this type and use it to represent these particular flowers; otherwise, a specialization of an existing model would represent them in the episodic situation.

As all of these frames become active, created, and/or revised, they become associated in various manners. First, their current forms, which may be specialized to varying extents, all become associated with one another. For example, the specialized frames for one's spouse and the flowers, together with the new frame for the vase and the unspecialized frame for the table, all become inter-associated. Second, the current forms of these frames become associated with the living room in the world model and with the temporal index *just before dinner* on this particular day. Thus, the episodic contributions to this episodic situation are specializations of the frames for one's spouse and the flowers, the new frame for the vase, and the associations between the individuals, models, location, and time. The generic contributions include the non-specialized aspects of all frames currently active, along with the location in the world model and the active temporal structures.

In this example, we have thus far ignored the fact that an episodic situation contains a *series* of images, not just one. Assuming that the individuals and location for a situation remain constant, several additional factors must be added to capture changes over time: First, the positions of individuals in the situation must be represented to capture their changes in position from image to image. We assume that positional information arises, first, from qualitative distinctions in the configuration of perceptual symbols *within* an image (e.g., the vase is on/not on the table), and second, from qualitative differences in configurations *between* images (e.g., one's spouse is closer/farther/same distance from the table). Second, the specializations of individuals may vary from image to image, as when a vase on a table in one image falls to the floor and breaks in subsequent images (i.e., the vase's state changes from *intact* to *shattered*, as depicted by perceptual symbols in different images). Finally, the images in an episodic situation are associated serially, as are the corresponding individuals within them, to capture temporal sequences in the physical event.

Generic Situations

Generic situations capture important constancies across related episodic situations. Like models, generic situations do not have direct counterparts in the physical world but have the ontological status of mental types. Two key issues are: What determines when episodic situations are related? What determines the important constancies abstracted into a generic situation? We propose a single answer to both questions. Two situations are related when:

(1) They share a common number of discrete informative images.
(2) They share common individuals and/or models.
(3) The configuration of individuals/models in each analogous image between situations is qualitatively the same.
(4) The transformations of individuals/models between analogous images is qualitatively the same.
(5) The two situations culminate in a common conclusion state.

Once these commonalties have indicated that two episodic situations are related, they are abstracted into a generic situation for this type of situation.

Establishing commonalties. When identifying and forming a generic situation, it is essential to determine which individuals and models constitute important commonalties, because a pair of situations could present many possibilities. A variety of factors probably constrains this process, including:

(1) Individuals not common to multiple episodic situations are likely to be excluded from the generic situation.
(2) Individuals not common across all images within a single episodic situation are likely to be excluded.
(3) If an individual is not common across all images within a single episodic situation but is common in the analogous images of two episodic situations, it should be considered.

These three 'syntactic' criteria may provide considerable mileage in determining, first, that two episodic situations are related, and second, which individuals should be extracted. Two, more conceptual, factors are undoubtedly important as well:

(4) When background beliefs specify that an individual or model should be relevant to a generic situation, it should be considered.

(5) When background beliefs specify that an individual or model should not be relevant to a generic situation, it should be dropped.

Consider acquiring the generic situation for fueling a car at a filling station. Imagine that a jar occurs on the gas pump in every image of the first episode experienced, but doesn't occur at all in the second episode. Following the first principle, the jar would be excluded from the generic situation that covers this type of episodic situation. Alternatively, imagine that the first episode includes a bird sitting on the gas pump in the first image but not in any other. Following the second principle, the model for *bird* should not be included in the generic situation. Now imagine that the second episode also contains a bird on the gas pump, again only in the first image. Following the third principle, the model for *bird* might be included, although the fifth principle would probably rule it out, assuming that people generally believe non-human animals have little impact on mechanical processes (Murphy & Medin, 1985). In general, however, the fifth principle might typically not need to be invoked, because of the low likelihood that a bird will occur in the first image of every subsequent episode. In contrast, if the first image of both episodic situations contains an adult human who pulled a lever on the gas pump and then was absent in all subsequent images, the third principle would cause *adult human* to be abstracted. Because the first image of all subsequent events would probably also contain such an individual, it would remain in the generic situation. This individual might also remain if background knowledge specifies that humans can play critical roles in mechanical processes (i.e., the fourth principle).

An alternative rule for identifying related generic situations might be that multiple episodes must simply share a common concluding state, rather than having to share common configurations of individuals and models in prior states as well. Certainly, this is an empirical issue, but our intuition is that separate series of configurations leading to a common concluding state constitute two different generic situations. To see this, imagine that a new technology develops for fueling cars, where the driver removes an empty gas tank from a car and replaces it with a new tank full of gas, using a mechanical cart to roll the tanks in and out of the car at a filling station. This new technology achieves the same concluding state as the old technology: a car with a full tank of gas. However, we suspect that people would represent the two means of achieving this state as different generic situations, because the individuals differ qualitatively in prior images.

We further suspect that common *configurations* of individuals are important in determining generic situations, not just common individuals. Consider the difference between self service and full service at gas stations. Both contain the same individuals but in different configurations. In self service, one operates the pump and pays the attendant from outside the car, whereas in full service, the attendant operates the pump and collects payment at a car window. Again, our intuition is that these constitute different generic situations, because of their different configurations of individuals.

Levels of abstraction and transcendence. Generic situations vary considerably in their level of abstraction. At one extreme, related episodic situations contain exactly the same individuals, at the same location in a world model, at the same relative (not absolute) time. For example, if someone only puts gas in her car at the same pump, at the same gas station, in the evening, then a highly specific generic situation develops, containing the same individuals, at the same location in the world model, at the same relative time. At the other extreme, related episodic situations contain no identical individuals, do not occur at the same location in a world model, and do not occur at the same relative time. For example, if someone has observed different people gassing their cars at different gas stations at different times of day, then a relatively abstract generic situation develops, containing no individuals but only models, at different locations in the world model, at different times. In between lie many intermediate levels of abstraction. For example, someone could always fill her car at the same gas station at the same time of day but at different pumps; or someone could fill his car at different gas stations but always at the same relative time; etc. Also, models at different levels of abstraction may reside in different generic situations. For example, if someone only puts premium gas in a car, the model for *premium gas* resides in the generic situation, not *gas;* similarly, if one only ever puts gas in cars and not in trucks, motorcycles, etc., then the model for *car* resides in the generic situation, not *vehicle.*

To the extent that a generic situation occurs across multiple locations, it transcends the world model. If one buys gas at a wide variety of gas stations, the generic situation for gassing one's car becomes transcendent. However, if a certain location has a much higher strength of association to the generic situation, it may tend to ground the generic situation there, at least to some extent.

As the abstractness of a generic situation increases, transcendence of individuals may also occur, although in a somewhat different sense from transcendence of world models. The basic idea is that abstract generic situations contain abstract models of individuals rather than frames for individuals per se, thereby transcending particular individuals in the world (e.g., particular gas station attendants). The result is a transportable generic situation, widely applicable to a variety of physical situations, not just one narrow type.

Constructing generic situations. We propose that two mechanisms discussed earlier for the construction of models--superordinate aspects and episodic inter-association--are also central to the construction of generic situations. Once commonalties between two related episodes have been identified, superordinate aspects are established within the generic situation as follows:

(1) If an individual frame represents an individual common to both episodic situations, then specializations of this frame from the two episodic situations, together with their superordinate aspects (if the specializations differ), are added to the generic situation.

(2) If a model frame represents an individual common to both episodic situations, then specializations of this frame from the two episodic situations, together with their superordinate aspects (if the specializations differ), are added to the generic situation.

(3) If a model frame represents two different individuals across the episodic situations, then specializations of this frame for the two individuals, together with their superordinate aspects, are added to the generic situation.

(4) Superordinate aspects for all frames in the generic situation are set to specializations extracted from the second episodic situation, while specializations from the first episodic situation remain associated.

Additionally, the generic situation becomes episodically organized as follows:

(5) The current forms of all frames for individuals and models are associated to one another.

(6) The current forms of all frames for individuals and models are associated to the world model locations of the two episodic situations.

(7) The current forms of all frames for individuals and models are associated to temporal indices for the two episodic situations.

To see how this might work, again consider learning to gas a car, with the first two episodic situations occurring at the same gas station in the evening but using different pumps to fill the same car. On experiencing the second situation, a generic situation is established, because the five criteria for relatedness are satisfied. Because the repeated car is a significant individual, it has an individual frame, which is included in the generic situation. If its tank was half full on the first occasion and near empty on the second, a superordinate aspect for *gas level* is established in the individual frame (if not present already) and specialized to *near empty*. All other aspects of the car extracted from these two situations are processed similarly.

Although the pay booth remains constant in both episodic situations, it may not have an individual frame, assuming it is not significant. Consequently, the model for pay booth is inserted into the generic situation, specialized the same way on each occasion, unless different aspects of the booth were noted. Because the pumps differ on the two occasions, the model for *gas pump* is inserted into the generic situation, specialized to reflect any differences in the pumps that were noted, with superordinate aspects for specializations also being computed as needed. For both the pay booth and pump, the most recent specializations of superordinate aspects reside in the current form of the model. Finally, the current forms of all frames for these individuals and models are inter-associated, together with the world model location and relative time, which remain constant.

Once a generic situation has been established, new episodic situations related to it are integrated according to the following principles:

(1) If any specializations in the new episodic situation aren't covered by existing superordinate aspects in the generic situation, new superordinate aspects are computed to cover them.

(2) The current specializations in all frames for individuals and models are set to those in the current episodic situation, while those from all previous episodic situations remain associated.

(3) The current forms of all frames in the generic situation are inter-associated, together with the current world model location and time.

For example, if a third episodic situation of pumping gas occurs at a new gas station at a new time, superordinate aspects in most frames, except those involving the car, will probably need to be computed, assuming that their specializations were extracted and differ. Additionally, the generic event will now be associated to a new world model location and time, while remaining associated to the older ones, thereby beginning the process of making the generic situation transcendent.

Specialized frames in generic situations. As a consequence of the procedure just described, the frames for individuals and models in a generic situation become tailored to the episodic situations that instantiate it. Because superordinate aspects in a frame are only computed over the episodes that instantiate a generic situation, and because these superordinate aspects, together with their specializations, are then associated with the generic situation, a specialized frame develops and becomes associated with the generic situation.

To see this, consider the superordinate aspects of *cloth* that might become established with the generic situation for getting gas. The cloths provided at most gas stations for checking oil are nearly always of a certain type, being about a foot square and made of a dark rough cotton. Because these individual cloths differ from station to station, a model represents them in the generic situation; but because these cloths tend to be similar, and because they tend to differ from many other cloths, the model for *cloth* in this generic situation is specialized to the episodic situations that instantiate it. Similarly, cloths in other generic situations are tailored to their episodic situations, such as the cloths for changing a baby's diapers, the cloths for cleaning glasses, and so forth. Different specializations of *cloth* develop in these other generic situations, having different superordinate aspects and specializations.

In this manner, specialized frames for models develop locally in generic situations. Rather than there being a single universal model that represents a type of individual across all situations, specialized models represent the particular kinds of individuals in specific situations. Analogously, specializations of individuals develop across generic situations, such as different specializations of the same person at work versus parties.

Importantly, specialized frames for models and individuals continue to obey the one-entity one-frame principle. Although specialized models develop in different situations, they nevertheless reside in a single frame, because they all represent the same type of individual. For example, specialized models of *car* in the generic situations for buying a car, driving a car, and getting gas constitute different perspectives on the same type of individual, namely, *cars*. Consequently, information about cars across all situations is integrated into a single frame. Specialized models develop as subsets of information in the frame become associated with particular situations. In a particular situation, the superordinate aspects and specializations associated previously with the situation become active to produce a specialized model, while other information in the frame remains inactive.

Types of generic situations. Although there are many types of generic situations, several seem of particular interest. First are situations in which the parts of an entity are examined, such as listing the parts of a car (e.g., *wheels, door, bumper, hood, engine, fuel tank, seat, steering whee*l). Part listing situations may often develop through explicit instruction and formal education, although they may also result from having to search through the parts of an entity to achieve some goal (e.g., searching the parts of a car to find a malfunction).

Part listing might seem like the canonical or privileged situation that captures the 'true' concept for a category, namely, the concept that is universal and decontextualized. However, we suspect that listing parts is a situation essentially the same as all others. One reason is that people appear to adopt perspectives while listing the parts of an entity. For example, people may begin listing the parts of a car by viewing it initially from the side (e.g., *wheels, door, window*). As listing continues, however, the perspective changes, perhaps viewing the car from the front (e.g., *bumper, hood*), then from under the hood (e.g., *engine, fuel tank*), and then from the passenger area (e.g., *seat, steering wheel, glove compartment*). Such shifting of perspective would suggest that people are imagining a physical situation in which they view a car and describe it. A second reason for believing that listing parts constitutes a situation is that we suspect people often imagine contextual information, such as pavement below a car or water surrounding a fish. Finally, we suspect that the models in part listing situations are specialized for these particular situations. Many more parts of a model may become associated with a generic situation for part listing than for situations in which part listing is not central. People's most extensive inventories of parts for a model may reside in these situations. In other situations, only the minimal parts needed to identify the model, as well as parts relevant to the situation, are represented. For example, only the major exterior parts of a car may be represented in the situation for getting gas, along with the fuel tank and fuel tank cap, because of their relevance to this situation.

Two predictions arise from this view: First, subjects should list the most parts in the part listing situation and list fewer parts in other situations. Second, minor parts relevant to a goal-directed situation may be omitted inadvertently from part listing situations, because they are not perceptually salient. In the part listing situation for *car*, people may not represent *gas tank cap*, because of its low perceptual salience in this context. In contrast, people would probably produce this part in the situation for getting gas, because it is so relevant.

Several other types of situations also appear important. First, we suspect that goal-directed situations are relevant for many models. In these situations, agents achieve goals, often with respect to objects, using instruments, and so forth. A given model may be an agent in some situations, an object in others, an instrument, recipient, or any of many other possible roles. For example, *dog* is an agent in the situation for chasing a cat, an object in the situation for being wormed, and an instrument in bird hunting. For every goal to which a model is relevant, there is a corresponding situation that contains a specialization of it.

One final type of situation depicts origins. For many models, people know how individuals originate, such as mammalian reproduction for *dog,* seed germination for *tree,* and assembly line production for *car*. These situations, together with part listing situations that focus on internal parts and atomic/molecular structure, could be seen as constituting intuitive theories (e.g., Murphy & Medin, 1985). For example, the most definitive information for *dog* is arguably its reproductive origins and DNA structure. Situations that depict a dog's origins are those representing two dogs mating, a dog having puppies, and so forth. Situations that depict a dog's DNA structure might be classroom situations in which students see models of DNA. Interestingly, most people may use exactly the same DNA situation for all living things, because they can't represent the different DNA of each species. Consequently, these situations are not defining of species concepts, but are only a reference to linguistic expertise (Putnam, 1973).

Concepts Revisited

We began this section by reviewing the standard assumptions that concepts are context-independent and universal. According to this view, concepts are context-independent, because they are not represented in situations, and they are universal, because they represent all individuals of a concept across all situations. To the contrary, we have developed a framework in which models are situated and local. Specializations of a model reside in situations and only cover individuals of the model that appear in that situation.

What is a concept on this view? We propose that a concept is the collection of all specialized models for a particular type of individual, together with their associated generic situations. The concept of *car*, for example, is the collection of all specialized models for *car*, together with their associated generic situations.[7]

What, if anything, brings coherence to the specialized models for a type? One might believe that a single general model ties all specializations across situations together. We are quite skeptical about this proposal, because we believe that the computation required to discover and maintain any information common to all specializations is beyond normal human capacity (Barsalou & Hale, 1993). There is also the related problem, well-known since Wittgenstein (1953), that many concepts may not possess information common to all instances across all situations. Most importantly, we don't believe that it is necessary for people to discover a common core to use concepts effectively. Knowing the appropriate use of a concept in each relevant situation is usually sufficient.

Nevertheless, people need to know that all of the situation-specific forms of a model are related. What mechanisms underlie this? We propose three: words, the one-entity one-frame principle, and conceptual chaining. Because the same word can refer to all specializations of a model, it organizes all of these specializations and indicates that they are related (e.g., "car" refers to cars in any situation). In the process, a word may well provide an illusion of semantic constancy, analogous to brightness and color constancy in perception (James, 1890/1950, p. 236). Although a constant word for a concept may create the illusion of a constant meaning, it actually ties together a diverse set of situation-specific models, with there being no guarantee of a common core.

The one-entity one-frame principle provides a second means of linking specialized models. Because all specialized models are manifestations of a single frame for a type of individual, they are linked together by a common frame. Specialized models arising out of the same frame are thus perceived as related to one another and different from specialized models arising out of different frames.

Finally, conceptual chaining provides a third means of linking specialized models. To see this, consider the following issue: Which model is used to represent an individual in an episodic situation that is unrelated to any previously experienced situation? Imagine that someone experiences a taxi cab for the first time. What specialized model of *car* represents the taxi in this episodic situation, and ultimately, in the corresponding generic situation? We propose that a specialized model is imported from the most similar episodic or generic situation. If taking a taxi for the first time is most similar to being a passenger in a personal car, then the specialized model from this generic situation is imported into the episodic situation for taking a taxi. Once imported, the model is specialized, perhaps adding yellow color, a meter, and a roof light. When another taxi episode occurs, the specialized model from the previous taxi episode can be imported into the second episode. If these two episodic situations are then perceived as related, a new generic situation is constructed to cover them. In this generic situation, the model for *taxi* is a specialized model that covers the two taxis in the two episodic situations. Much of this model's content was copied originally from the specialized model for *car* in the generic passenger situation. However, this model further contains unique perceptual symbols extracted from the two taxis, together with any superordinate aspects needed to integrate different specializations of the same aspect (e.g., two different types of roof light). The result is a specialization of *car* tailored to the cars in these related situations.

As this example illustrates, conceptual chaining links the specialized models for a given type across generic situations. In our example, the specialized model for *car* in the passenger generic situation led to the specialized model for *car* (i.e., *taxi*) in the taxi generic situation. If a new type of episodic situation were subsequently similar to the taxi generic situation, further chaining might occur. For example, the subsequent experience of auto couriers might produce a new generic situation, extending the model of *car* further by chaining the specialized model for *taxi* to a specialized model for *courier car*.

This sort of chaining process produces the radial structure of concepts noted by Lakoff (1987), producing heterarchical trees of models that project out from the most common specializations. We hasten to add, however, that these models reside in generic situations and are not decontextualized, similar to Lakoff's (1987) construct of idealized cognitive models. Again, there may be no single model that constitutes the core of the concept. Instead, the common word that refers to all specialized models provides a common thread, as do the common frame and the conceptual chains that link them together.

Further empirical evidence for conceptual chaining and radial categories comes from recent work by Malt on concepts such as *water* and *box* (Malt, 1991). In these studies, Malt finds that people use a term like "box" in referring to a diverse and somewhat incoherent set of referents, such that no core definition appears satisfied. Conceptual chaining, where specialized models of a concept are embedded in generic situations, provide a natural account of this phenomenon.

Preliminary evidence for generic situations in concepts. Clearly, empirical evidence is needed before we can accept the claims that concepts are situated and local. However, certain lines of research provide preliminary support. First, a variety of experiments have shown that instances of concepts are identified faster in context than in isolation. For example, Palmer (1975) and Biederman (1981) found that people identify visual objects faster in their expected locations than in isolation or in misleading contexts (e.g., a loaf of bread in a kitchen). Murphy & Wisniewski (1989) found that these situational effects are particularly strong for superordinate categories. Similarly, many investigators have found that people identify letters (i.e., concepts) faster in words (i.e., situations) than in isolation (e.g., Rumelhart & McClelland, 1982). All of these studies suggest that concepts are contextualized and not context-independent.

Situations are also implicated in the feature lists that subjects produce for concepts (Barsalou, Sewell, & Spindler, 1993). When subjects list the characteristics of a concept, such as *car,* they often mention aspects of the situations in which they occur, such as *garage, highway,* and *gas station.* The presence of such features again suggests that the concepts being described are situated.

Increasing work further suggests that concepts are local, not universal. When subjects judge the typicality of *birds* from the point of view of the average Chinese citizen, they produce very different graded structures than when they judge typicality from the American point of view (Barsalou et al., 1993; Roth & Shoben, 1983). For reviews of such findings, see Barsalou (1987, 1989).

Concepts and Meaning

In this final section, we address three common assumptions about meaning, argue that they are incorrect, and establish an account of meaning that we will compare with our account of concepts. We then address the relations between concepts and meaning, arguing that concepts are far from equivalent to meaning, but play important roles in establishing it. Finally, we extend our analyses of concepts and meaning to conceptual combination and metaphor.

Meaning

Our analysis rests heavily on the standard Fregean assumption that the meaning of an expression typically, but not always, contains both sense and reference, where a sense is a description that picks out individuals constituting the expression's reference (Frege, 1892/1952).[8] Together, the sense and the reference of an expression constitute its meaning. When addressing the relations between concepts and meaning, theorists sometimes make the following assumptions:

(1) Meanings are senses.
(2) Meanings are universal.
(3) Universal senses determine reference.

We address each assumption in turn and argue that all are incorrect.

Are meanings senses? Psychologists have traditionally equated meanings with senses, believing that understanding the concept associated with a word (i.e., its sense) explains its meaning. For example, the literatures on semantic memory and concepts have focused exclusively on the senses of words, ignoring reference (e.g., Smith, 1978; Rosch & Lloyd, 1978). Other approaches within the cognitive sciences often have this flavor as well, including Schank (1975, 1982), Langacker (1986), and Katz (1972). Critics within psychology have noted the absence of reference from psychological accounts of meaning and argued for its importance (e.g., Johnson-Laird, Herrmann, & Chaffin, 1984; Malt, 1991; Murphy, 1991). Malt (1991), for example, demonstrates that analysis of a word's extension can lead to conclusions about its sense that differ considerably from examining the sense in isolation.

Classic work in philosophy of language demonstrates the centrality of reference in meaning. Putnam (1973) argued that meanings "ain't in the head" but depend on referents in the environment. On his view, a term like "water" refers to a particular physical substance, regardless of people's beliefs about it (i.e., their concept or sense). Similarly, the *de re/de dicto* distinction illustrates the importance of reference in meaning (e.g., Kripke, 1977). Imagine that someone says, "Could you please bring me the mug on the table," when there is only a mug on the floor next to the table. The listener will probably assume that the meaning of "mug on the table" is actually the mug on the floor, such that the speaker's intended referent for "mug on the table" overrides its conceptual sense. In cases like this, the intuition is strong that the meaning of "mug on the table" is the actual mug on the floor and not a conceptualization of a mug on a table. Based on these arguments, and others to be presented shortly, reference appears central to meaning.

Are meanings universal? Cognitive scientists who study the relation between concepts and meaning typically address what we shall call *universal meaning*. A meaning is universal when it specifies the meaning of a word or expression with respect to the entire world. For example, the universal meaning of "bird," would be the conditions (i.e., the sense) that identify all birds in the universe (i.e., the extension) and reject all non-birds. Putnam's (1973) analyses of natural kind terms typically have this flavor, as when he tries to determine the conditions that pick out all instances of "water" in the universe and nothing else. Similarly, Malt's (1991) analyses of "water" and "box" have this flavor, as she attempts to specify family-resemblance senses that establish the universal extensions of these words (also see Murphy, 1991).

We certainly agree that meanings *can* be universal, as when theorists discuss the universal meanings of natural kind terms. However, we believe that universal meanings are simply a special case of *situation-specific meaning*. As we shall argue shortly, and as others have argued compellingly, meanings are *always* established with respect to a specific situation or discourse context (e.g., Clark, 1983; Clark & Clark, 1979; Nunberg, 1979). During normal communication, conversationalists continually establish and redefine common ground, such that words and expressions refer successfully to entities within it. If someone says, "I worked in my yard last weekend," entities in this situation are established as potential referents. If the speaker then says, "The lawn needed to be cut," the word "lawn" refers to the particular lawn in the speaker's yard. If the speaker had introduced a different situation by saying, "I worked at the park last weekend," and then said, "The lawn needed to be cut," the word "lawn" has a different reference. Most importantly, the reference in neither case is universal. "Lawn" does not pick out the entire set of lawns in existence (i.e., its extension); instead, it denotes two specific individuals in different situations.

We believe that universal meanings occur only when conversationalists establish the entire world or universe as common ground. Semantic theorists, for example, often make it clear that they are examining the meanings of words with respect to the entire set of entities in existence. Similarly, non-scientists on occasion may consider universal meanings, as when acquiring definitions for natural kinds in formal education. In general, however, we believe that universal meaning is relatively rare in normal conversation and cognition. Instead, we suspect that meaning is typically much more specific. People's poor ability to define words universally is consistent with this conjecture (e.g., Barsalou, et al., 1993; Malt, 1991; Rosch & Mervis, 1975).

Most importantly, we believe that excessive concern with universal meaning has distorted our understanding of the relations between meaning and concepts. When concepts are compared to this relatively rare and rarified variety of meaning, the typical relations between them are not apparent.

Do universal senses determine reference? Cognitive scientists often assume that universal senses associated with words determine their extensions. For example, Murphy (1991) argues that psychologists should focus on the study of intensions (i.e., senses), because once intensions are understood, they will determine extensions. Similarly, Malt (1991) attempts to discover the intensions that determine the complex extensions that people know for "water" and "box."

However, a variety of investigators have noted that expressions often establish reference *outside* their universal extensions. In these cases, universal senses do not determine reference, because they can only establish reference *inside* their universal extensions. For example, "ham sandwich" refers to a customer in a restaurant, not a sandwich, when a waiter says, "The ham sandwich requested more coffee" (Nunberg, 1979). Clearly, the universal sense of "ham sandwich" cannot be responsible for establishing reference in this sentence, because the universal sense can only establish reference to real ham sandwiches.

It is easy to find less exotic examples of reference to individuals that lie outside universal extensions (Barsalou, 1992a). If "Look at the horse" were uttered at an art museum, the referent might be a painting of a horse. If this sentence were uttered in a toy store, the referent might be a stuffed horse. If this sentence were uttered in a veterinary lab, the referent might be a vial of blood drawn from a horse. In each case, the referent lies outside the universal extension of real horses, such that the universal sense can't be establishing reference. As we shall argue shortly, other mechanisms, not universal senses, are typically responsible.

Relations between Concepts and Meaning

In assessing the relations between concepts and meaning, we will assume that: (1) meanings typically include reference, not just sense; (2) meanings are typically specific to particular individuals, not universal, and (3) meanings often extend beyond universal extensions.

Are meanings concepts? Recall that, on our view, a concept is a single frame representing a single type of individual, with this frame manifesting itself as different specialized models in different generic situations. Because the specialized models within a concept are tailored to situations, a concept is not a universal sense that represents properties true of all individuals. Instead, each specialized model in a concept represents properties true of some individuals in a specific situation.

This account of concepts differs significantly from meanings. First, consider perhaps the most typical case of meaning, where a word or expression refers to one individual (or a small set of individuals) in a particular situation (e.g., the examples for "lawn" earlier). In such cases, meaning includes both reference to an individual, as well as some sort of sense provided by the concept associated with the word or expression. As we suggest in the next section, these senses vary widely, often being much more specialized than entire concepts or universal senses. Thus, the standard case of meaning differs significantly from concepts, first, because it contains reference to an individual in a specific situation, and second, because its sense is tailored to the individual.

Next consider the relatively unusual case of universal meaning. When people use "car" universally, its meaning could potentially include both the concept for *car* as its sense, and the extension of all cars in the universe as its reference. Practically, however, we suspect that people are incapable of representing the entire concept and extension for "car" at a given point in time (Barsalou, 1987, 1989; Murphy, 1991). Instead, people only access a small subset of specialized models, and they only access a small subset of corresponding individuals in the extension. In either case, the universal meaning of "car" differs considerably from the concept for *car:* Whereas the universal meaning of "car" contains both sense and reference, the concept for *car* is just the sense.[9]

Finally, consider generics, another unusual variety of meaning. Generics may consitute the case where meanings most closely approximate concepts. When someone says, "I have no idea what a computer does," the reference of "a computer" is approximately the concept *computer.* When an expression refers to a concept in this manner, meaning and concepts appear roughly equivalent, because reference is to a concept, not an individual. Notably, the atypicality of generics as a variety of meaning underscores the non-equivalence of meaning and concepts.

Roles of Concepts in Meaning

Although concepts are typically not equivalent to meaning, they play three important roles in constructing it:

(1) Concepts establish reference,
(2) Concepts provide 'running commentary' about referents.
(3) Concepts establish domains of reference.

Whereas concepts function as senses in (1) and (2), they serve as referents in (3).

Concepts establish reference. We propose that the 'library' or 'chain' of generic situations that develop for a concept provides the basis of establishing reference from its associated word (Barsalou, 1992a). For example, the library of generic situations for *horse* enables people to establish the reference of "horse" in uses of this word across utterances.

Our account rests on two principles of establishing reference: *weak association* and *restricted discrimination*. By *weak association* we mean that a referent need only have a weak connection to a generic situation in a word's concept to become established as the word's meaning in some situation. Recall the sentence, "Look at the horse," for which "horse" could refer to a painting, toy, or blood sample, depending on the situation. Visual resemblance constitutes the weak association that enables references to the painting and toy. Actually being in the universal extension of "horse" is not necessary for being a referent. Instead, visual resemblance to a model for a horse in any generic situation is sufficient. For the blood sample, the fact that horses have blood constitutes another example of a weak association that enables reference. In this case, a part listing situation for *horse,* which includes *blood*, could provide the connection, although people may have to use a part listing situation for *mammal*, if their part listing situation for *horse* does not include *blood*.[10]

Because weak association is often sufficient to establish meaning, a potential problem is that too many entities in the world could be the referent of a word. For "Look at the horse," any entity in the world that could be connected to *horse's* library of generic situations in some weak manner is a potential referent. *Restricted discrimination* typically eliminates this problem. Over the course of a conversation, speakers establish and adjust the domain of potential reference (i.e., common ground; Clark & Marshall, 1981; Clark & Schaefer, 1989). At any given point, only a small subset of entities in the world is under discussion, and these entities provide the most likely referents of an utterance. Certainly, the set of potential referents may change, but at any point, only a small subset is typically relevant. When someone at a museum says, "Look at the horse," if the speaker and listener have established that the domain of reference is the current room, then "horse" probably refers to something in the room, not to any possible thing in the world. Because the domain of reference is limited, the intended referent experiences little or no competition from other possible referents outside the domain. For example, there is no competition from horses or anything horse-like outside the museum, because these entities lie beyond the current domain of reference. Consequently, reference can easily be established to a picture of a horse, assuming that there is only one such picture in the room.

Moreover, weak association enables 'creative' reference, once a limited domain of reference has been specified. It becomes possible to establish reference to something that isn't actually a horse but only a painting of one, because the painting is the closest thing to a horse in the current domain of reference. In this way, weak association extends a word's meaning in creative ways beyond its universal extension, while restricted discrimination allows listeners to identify the new meaning. Together, these two mechanisms produce frequent exceptions to the principle that universal senses establish reference.

These two mechanisms, together with our definition of *concept,* distinguish literal from figurative meaning. We define *literal meaning* as reference by a word to individuals *within the type* of its associated concept. "Horse" is used literally, for example, whenever it refers to individuals of the conceptual type *horse* (i.e., real horses). Whenever "horse" refers to individuals outside this type--to anything but a real horse--it is used figuratively. For example, using "horse" in referring to a picture of a horse, a toy horse, or the blood of a horse constitute cases of metonymy, not literal usage (Gibbs, in press;

Lakoff, 1987). In these cases, the joint action of weak association and restrictive discrimination establish reference outside the conceptual type.

Finally, weak association and restricted discrimination make universal meaning irrelevant to most normal conversation. Because reference is typically restricted to specific individuals within specific situations, weak association and restrictive discrimination are usually sufficient to establish reference--universal senses that determine universal extensions are unnecessary. Because most people never have to discriminate the universal extension of a word from its complement, they never acquire universal senses, nor can they provide them.

Concepts provide 'running commentary.' Thus far, we have assumed that the only role of concepts in meaning is to establish reference. Following Frege again, however, concepts also serve to describe referents, providing a 'running commentary' on them. In Frege's well known example of "morning star" and "evening star," the referent remains constant, while the running commentary on it diverges. Not only does the concept for each of these two expressions establish reference, it also draws attention to the generic situation in which Venus is currently perceived, either the morning sky or evening sky.

We propose that running commentary often draws on the specialized frames for individuals and models in generic situations. Once an entity is established as a referent in a particular situation, a specialized frame from the corresponding generic situation provides running commentary on the referent. To see this, consider how the senses of "Bill Clinton" vary when this name refers to Clinton in different situations, such as speaking to the press, jogging through a Little Rock neighborhood, or joining a jazz band on stage. In the press situation, a specialized model might represent Clinton as upbeat but evasive; in the jogging situation, a specialized model might represent Clinton as slightly overweight and hungering for a hamburger; in the jazz situation, a specialized model might represent Clinton as an earnest but mediocre saxophonist. Although the same name, "Bill Clinton," establishes the same reference in each situation, the running commentary on the referent varies, depending on the specialized model that the current situation activates. In this manner, frames for individuals and models may provide the basis of running commentary. Running commentary can further arise from knowledge about referents and through metaphor, as we argue in the section on Extensions.

Concepts establish domains of reference. Generic situations establish domains of reference when speakers introduce them in a conversation, thereby designating entities within them as potential referents. If someone says, "There's something you should know about gassing your car," she establishes a generic situation that becomes the current domain of reference. If she then says, "Most gas stations check your air when you request full service," listeners can easily compute the references for "air" and "full service"--using weak association and restricted discrimination--as entities and events within the generic situation.

Besides providing potential referents for utterances, generic situations also play central roles in constructing new situations and in modifying existing ones. For example, if someone says, "When you get gas, it's also possible to buy lottery tickets," the generic situation for buying lottery tickets is embedded in the generic situation for buying gas. Once these complex situations are established, entities within them become potential referents, as in "I bought a ticket and some oil."

As new discourse situations develop, they feed back into the libraries of generic situations for relevant words. Following a basic law of memory, once information is processed extensively in working memory, it becomes established in long-term memory. On hearing that lottery tickets can be purchased together with gas, a listener might insert a specialized situation for buying lottery tickets into the generic situation for getting gas.

Extensions

Conceptual combination. One way of viewing conceptual combination is as the process of forming new concepts from old ones. Our view is that people rarely establish new concepts during conceptual combination but primarily specialize existing ones, as in the Smith, Osherson, Rips, and Keane (1988) modification model (also see Cohen & Murphy, 1984; Murphy, 1988, 1990). For example, people do not establish new concepts for *large table* and *blue fish*, but specialize existing concepts for *table* and *fish.* Following the one-entity one-frame principle, people do not construct new frames, because *large table* and *blue fish* do not constitute significant new types of entities. Certainly, some conceptual combinations may constitute significant new types and therefore warrant new frames (e.g., *pet fish* for someone newly interested in aquariums). In general, though, most conceptual combinations may simply modify existing frames.

Another frequent assumption about conceptual combination is that the meaning of a combination follows directly from the composition of its constituent concepts. Our view is that this may happen occasionally, but that, at least as often, the meaning of a combination is determined through reference to a situation (Downing, 1977), or what Hampton (1988) calls *extensional feedback* (also see Clark & Gerrig, 1983; Cohen & Murphy, 1984; Gerrig & Murphy, 1992). If you know that there is a house on Mary's block with a big smile painted on the front, and if I say, "Mary drove by the happy house," you will probably assume that the meaning of "happy house" is this particular house. Additionally, you would probably assume that this house is 'happy' is because of its 'smile.' Thus, a property of the referent--its 'smile'--not only establishes reference to this particular house, it also determines the running commentary associated with the conceptual combination. Alternatively, if I had mentioned instead that a joyous event had occurred recently at this house, the referent would remain the same, but the running commentary would differ, with the house now being 'happy' because of the joyous event. In these examples, the sense of "happy house" does not result from simply combining the concepts for *happy* and *house;* instead, the sense depends on knowledge of the referent (Murphy, 1988).

Consideration of *blue sand* and *blue car* further demonstrates the importance of reference in conceptual combination (Barsalou, 1993). Although *blue* is the modifier in both cases, it modifies each noun differently, depending on what is known about the referent's color. Because mass nouns tend to have a color throughout their mass, "blue sand" is conceived as blue throughout. In contrast, because count nouns often tend to have different surface and interior colors, with the surface color being more salient, "blue car" is conceived as only blue on its body. In each case, the sense of the conceptual combination is established through reference, rather than through a purely conceptual combination of senses. The meaning of the combination does not simply arise conceptually from combining the concept for *blue* with the concept for the head noun, as in adding a feature for *blue* to the feature list for the noun, or as in modifying the slot for *color* in the noun frame with *blue.* Instead, the different referents inform the sense of *blue*, depending on their typical coloration.

Because conceptual combinations are generally produced in discourse contexts, they refer to entities within situations, with reference often playing a central role in establishing sense. Studying conceptual combinations in isolation, out of context, fails to provide an accurate account of how they typically behave, because the referential component is minimized (although it is extremely likely that subjects establish reference situations in an uncontrolled manner to make sense of the task). Because the meaning of a conceptual combination typically depends on the situation in which it occurs, it is important, if not essential, to study conceptual combination under these conditions.

Metaphor. Recently, Lakoff and others have argued that experiential gestalts provide metaphorical understandings of abstract concepts (Gibbs, in press; Lakoff, 1987; Lakoff & Johnson, 1980).[11] For example, the experiential gestalt of heated liquid exploding out of a container provides a metaphorical understanding of anger, as in "John blew his stack." We agree that such metaphors provide important understandings of entities and events, as we discuss shortly. However, we do not believe that experiential gestalts constitute the most basic meanings of many metaphorical expressions. Instead, we believe that the primary basis of these meanings lies in direct experience. If someone believes that the meaning of "blow one's stack" is the experiential gestalt of an exploding container, he does not possess an adequate understanding of this expression. To understand it fully, one must have direct knowledge of the introspective states associated with anger, as well as the behavioral actions, the initiating situations, and so forth (Stein & Levine, 1990; Stein et al., in press). A full understanding of "blow one's stack" requires experiential knowledge of the actual situation--knowledge of a metaphorically related domain is not sufficient.[12]

Situations do exist in which an utterance can only be understood metaphorically. When people lack any experiential knowledge of a domain, the meaning of a sentence about it must be established through metaphor. Consider a man's understanding of a woman saying to him, "During labor, my contractions were like an endless string of hamstring cramps." Because the man has never experienced labor contractions, but has probably experienced hamstring cramps, his sole means of understanding "contractions" is the source domain, hamstring cramps. Under such conditions, metaphors may provide the only source of meaning. When people have direct experience of the target domain, however, the direct experience also becomes an important source. For example, women who have experienced labor possess direct knowledge of contractions, such that "contractions" in the above sentence is understood primarily in terms of experiencing contractions, with "hamstring cramps" representing the limited understanding that the man, to whom the utterance was directed, could have of them.

As this example illustrates, the meaning of a metaphorically motivated utterance is represented at several levels. At the most basic level, the meaning of such utterances rests on their reference. If someone says, "Were you surprised at the meeting yesterday when I blew my stack," the most basic part of the meaning for "blew my stack" is the actual event to which the speaker refers. Whatever memory speakers have of this event is the referent of "blew my stack" and constitutes its most basic source of meaning.

However, other levels of conceptualization enter into the meaning as well. Both the speaker and listener possess generic situations for being angry, which are represented in terms of perceptual symbols for this type of event. These perceptual symbols may have been extracted from the introspection of anger, the visual and auditory perception of another person's angry behavior, and so forth. Most importantly, this generic situation is built up from direct experience for this type of event--it is not metaphorical. Consequently, when someone says he blew his stack, the referent event is conceptualized, at least partly, in terms of the generic situation developed from this type of experience. In a Fregean manner, the generic situation provides a sense, or running commentary, on the referent event.

In the absence of a clear referent, these experientially-based generic situations may provide possible referents for meaning. If I say, "I blew my stack this morning on discovering that the newspaper was soggy," you have no direct experience of this event. Consequently, you must use your knowledge of generic situations to establish a domain of reference, creating an imagined episodic situation by combining generic situations for picking up newspapers and being angry. Once established, these generic situations provide direct experiential referents for "blew my stack." You may specialize them if you know something about where I pick up my newspaper, and if you are familiar with my expressive mannerisms. Even in such cases, though, the basic meaning for "blew my stack" is grounded in direct experience, not in a metaphorically related domain.

We hasten to add, however, that metaphorical conceptualization may provide important running commentary on referent events. For example, an image of liquid exploding out of a container may well exist simultaneously with images of the referent event and the corresponding generic situation. Most importantly, it may play central roles in conceptualizing the referent event. First, it may highlight a particular aspect of the referent event and make the corresponding generic situation salient. For example, "blew my stack" may focus attention on the particular moment when anger was released, as opposed to some other part of the event (e.g., the initiating conditions, ensuing actions). Subsequently, the generic situation that applies to the release of anger may become part of the sense, and not other generic situations that apply to other aspects of the event.

Second, the metaphorical level may contribute significantly to beliefs and understandings about the nature of the referent event (Gibbs, in press; Lakoff, 1987; Lakoff & Johnson, 1980). For example, the exploding container metaphor may lead speakers to believe that emotions reflect psychic energies in 'regions of the mind,' breaking forth from time to time to produce behavioral outbursts. In this way, the metaphorical level may provide intuitive theories about aspects of direct experience that remain unobservable (Wellman & Gelman, 1988). We agree that these metaphorical understandings may be the most central aspects of metaphorical expressions on many occasions.[13]

Finally, the metaphorical level may provide an entertaining running commentary on the referent event. It is amusing to imagine steam spewing out of a person's ears as she releases her anger. Again, however, such conceptualizations do not constitute the most basic meaning of an utterance, because they do not constitute the referent events. Although the metaphorical level may play central roles in meaning, it usually accompanies referent events and generic situations that are grounded in direct experience.

Conclusion

Again, we stress that this working paper represents a theory in the early stages of development. Our theory clearly requires considerably more empirical support, as well as more precise articulation. However, the primary goal of this paper has been to outline our theory in its current form so that we can begin to examine its claims empirically and implement it computationally. Empirical and simulation projects currently underway will hopefully increase our understanding of these issues, producing a more refined and sophisticated theory in the process.

Acknowledgment

For helpful comments on this paper, we are grateful to Raymond Gibbs, Barbara Malt, and Gregory Murphy, who do not necessarily endorse our position in its entirety. The first section on perceptual symbols is, to some extent, a summary of Barsalou (1993). The remaining three sections are the result of weekly discussions over the past year initiated by the last five authors, whose order of authorship was determined randomly. Address correspondence to Lawrence W. Barsalou, Department of Psychology, 5848 S. University Ave., Chicago, IL 60637.

Footnotes

[1] We will use italics when referring to conceptual or semantic content, and double quotes when referring to linguistic forms. Thus, *apple* refers to a concept or meaning, whereas "apple" refers to the respective word. We shall also use italics at times for emphasis, with the usage being clear from context.

[2] It is important to note that "proposition" has a variety of senses, including (1) the underlying conceptual gist of a sentence or scene, (2) the representation of this gist using abstract, amodal, and arbitrary languages, such as propositional logic or predicate calculus, and (3) a specific type of language for representing gist, namely, propositional logic. Throughout this paper, we primarily use the second sense. As shall be seen, however, propositions in the first sense can be expressed though the perceptual symbols discussed shortly, not only through the 'languages of thought' associated with the second sense.

[3] These images are not compositional in the strong sense of being exactly the sum of their parts with no emergent properties.

[4] Further power for representing abstract concepts arises through metaphorical mappings from perceptual domains (e.g., Lakoff & Johnson, 1980). However, we believe that most abstract concepts also have direct perceptual bases in experience, as discussed later.

[5] Medin and Schaffer (1978) and Nosofsky (1988) present a different view of *exemplar* that is more like the construct of *individual* presented shortly.

[6] In this and all remaining examples, we focus on visual information. As noted earlier, however, it is essential to remember that images, situations, and events typically contain perceptual symbols from all aspects of experience, including introspection and proprioception.

[7] The first author has been fairly schizophrenic in his definition of *concept*.. Barsalou (1987, 1989, 1993) argues that concepts are temporary representations in working memory; Barsalou (1992b) and this chapter argue that concepts are bodies of knowledge in long-term memory; Barsalou (1992a) argues that concepts determine categorization. From this paper on, the following terminology will be followed: *Concepts* are the underlying knowledge in long-term memory from which temporary *conceptualizations* in working memory are constructed.

[8] Contrary to standard Fregean analysis, we shall assume that senses are psychological entities, rather than ideal descriptions that exist independently of human observers (Barsalou, 1992a).

[9] The question arises as to whether concepts have reference. To examine this issue, consider the specialized model for *car* in the generic situation for getting gas. One possibility is that this specialized model has as its referents all individuals that led to its construction. Alternatively, its referents might be all individuals to which it potentially could be applied, past, present, and future. Our current position is that inactive concepts have *potential reference* but not *engaged reference*. For example, a specialized model of *car* could be used to represent a particular car in a past episode at a gas station, to represent a particular car currently being filled at a gas station, and so forth. Although specialized models can be engaged in reference, as these examples illustrate, they usually lie dormant and unengaged, not having any current reference, only having potential for reference.

[10] Some weak associations may establish reference more naturally than others. For example, reference from "horse" to a stuffed horse may be easier than to a saddle, even though the saddle is clearly associated. Certain references may be awkward or impossible for a variety of reasons (Nunberg, 1979), indicating that limiting conditions on weak association must be established. Additionally, *episodic situations*, not just generic situations, may serve as the basis of weak association, as in Clark and Clark's (1979) example of "to Rockefellar the crowd."

[11] Following Lakoff, our use of *metaphor* in this section does not refer in the narrow sense only to metaphorical tropes. Instead, we employ *metaphor* in the looser sense of using knowledge from one domain to understand another.

[12] Interestingly, our view is entirely consistent with the central assumption of cognitive linguistics that semantics is experiential and embodied. Where our view diverges is that cognitive linguists have not stressed that abstract domains may be understood through direct experience, but have stressed instead that abstract domains are understood in terms of other 'more accessible' domains. Our disagreement is essentially that abstract domains are actually quite accessible and that they are understood to a large extent in terms of direct perceptual experience.

[13] For some metaphorical expressions, though, the establishment of their reference may be so habitual and automatic that the metaphorical level no longer becomes active, much like the literal meanings of highly familiar indirect speech acts (e.g., Could you pass the salt; Gibbs, in

press). For example, someone may use the expression, "hit the hay," so often that it is only conceptualized as getting into bed, no longer accompanied by a metaphorical conceptualization of sleeping in a barn.

References

Anderson, J.R., & Bower, G.H. (1973). *Human associative memory.* Washington, DC: Winston.
Barsalou, L.W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 629-654.
Barsalou, L.W. (1987). The instability of graded structure in concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (101-140). New York: Cambridge University Press.
Barsalou, L.W. (1988). The content and organization of autobiographical memories. In U. Neisser & E. Winograd (Eds.), *Remembering reconsidered: Ecological and traditional approaches to the study of memory* (193-243). New York: Cambridge University Press.
Barsalou, L.W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (76-121). New York: Cambridge University Press.
Barsalou, L.W. (1991). Deriving categories to achieve goals. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, 1-64). New York: Academic Press.
Barsalou, L.W. (1992a). *Cognitive psychology: An overview for cognitive scientists.* Hillsdale, NJ: Erlbaum.
Barsalou, L.W. (1992b). Frames, concepts, and conceptual fields. In E. Kittay & A. Lehrer (Eds.), *Frames, fields, and contrasts: New essays in semantic and lexical organization* (21-74). Hillsdale, NJ: Erlbaum.
Barsalou, L.W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, & M.A. Conway (Eds.), *Theories of memories* (29-101). London: Erlbaum.
Barsalou, L.W., & Hale, C.R. (1993). Components of conceptual representation: From feature lists to recursive frames. In I. Van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (97-144). San Diego, CA: Academic Press.
Barsalou, L.W., Sewell, D.R., & Spindler, J.L. (1993). *Flexibility and stability in concepts.* Manuscript in preparation.
Bickhard, M.H. (1980). *Cognition, convention, and communication.* New York: Praeger.
Biederman, I. (1981). On the semantics of a glance at a scene. In M. Kubovy & J.R. Pomerantz (Eds.), *Perceptual organization* (213-253). Hillsdale, NJ: Lawrence Erlbaum Associates.
Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115-147.
Carlson, G.N. (1982). Generic terms and generic sentences. *Journal of Philosophical Logic, 11,* 145-181.
Church, A. (1956). *The problem of universals.* Notre Dame, IN: University of Notre Dame Press.
Clark, E.V., & Clark, H.H. (1979). When nouns surface as verbs. *Language, 55,* 767-811.
Clark, H.H. (1983). Making sense of nonce sense. In G.B. Flores d'Arcais & R. Jarvella (Eds.), *The process of understanding language* (297-331). New York: John Wiley.
Clark, H.H., & Gerrig, R.J. (1983). Understanding old words with new meanings. *Journal of Verbal Learning and Verbal Behavior, 22,* 591-608.
Clark, H.H., & Marshall, C. (1981). Definite reference and mutual knowledge. In A. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding* (10-63). New York: Cambridge University Press.
Clark, H.H., & Schaefer, E.F. (1989). Contributing to discourse. *Cognitive Science, 13,* 259-294.
Cohen, B., & Murphy, G.L. (1984). Models of concepts. *Cognitive Science, 8,* 27-58.
Conway, M.A. (1990). *Autobiographical memory: An introduction.* Buckingham, UK: Open University Press.
Downing, P. (1977). On the creation and use of English compound nouns. *Language, 53,* 810-842.
Fauconnier, G. (1985). *Mental spaces.* Cambridge, MA: MIT Press.

Finke, R.A. (1989). *Principles of mental imagery.* Cambridge, MA: MIT Press.
Fodor, J.A., & Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28,* 3-71.
Frege, G. (1952). On sense and reference. In P.T. Geach & M. Black (Eds. and Trans.), *Philosophical writings of Gottlob Frege.* Oxford: Basil Blackwell. (Original work published 1892.)
Garner, W.R. (1974). *The processing of information and structure.* New York: Wiley.
Garner, W.R. (1978). Aspects of a stimulus: Features, dimensions, and configurations. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and categorization.* Hillsdale, NJ: Erlbaum.
Gentner, D. (1989). The mechanisms of analogical reasoning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (199-241). New York: Cambridge University Press.
Gernsbacher, M.A., Varner, K.R., & Faust, M.E. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 16,* 430-445.
Gerrig, R.J., & Murphy, G.L. (1992). Contextual influences on the comprehension of complex concepts. *Language and cognitive processes, 7*, 205-230.
Gibbs, R.W., Jr. (in press). *Poetics of mind: Figurative thought, language, and understanding.* New York: Cambridge.
Glenberg, A.M., & Langston, W.E. (1992). Comprehension of illustrated text: Pictures help to build mental models. *Journal of Memory and Language, 31,* 129-151.
Glenberg, A.M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Language, 26,* 69-83.
Goldin-Meadow, S., Alibali, M.W., & Church, R.B. (1993). Transitions in concept acquisition: Using the hand to read the mind. *Psychological Review, 100*, 279-297.
Goldstone, R.L. (in press). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
Goldstone, R.L., & Medin, D.L. (in press). The time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
Hampton, J.A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 12-32.
Harnad, S. (1987). Category induction and representation. In S. Harnad (Ed.), *Categorical perception: The groundwork of cognition* (535-565). New York: Cambridge University Press.
Harnad, S. (1990). The symbol grounding problem. *Physica D, 42,* 335-346.
Huttenlocher, J. (1973). Language and thought. In G.A. Miller (Ed.), *Communication, language, and meaning: Psychological perspectives* (172-184). New York: Basic Books.
Huttenlocher, J. (1976). Language and intelligence. In L.B. Resnick (Ed.), *The nature of intelligence* (261-281). Hillsdale, NJ: Erlbaum.
Huttenlocher, J., & Higgins, E.T. (1978). Issues in the study of symbolic development. In W. Andrew Collins (Ed.), *Minnesota Symposia on child psychology,* (Vol. II, 98-140). Hillsdale, NJ: Erlbaum.
James, E. (1950). *The principles of psychology* (Vol. 1). New York: Dover. (Original work published 1890.)
Johnson-Laird, P.N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.
Johnson-Laird, P.N., Herrmann, D., & Chaffin, R. (1984). Only connections: A critique of semantic networks. *Psychological Bulletin, 96*, 292-315.
Katz, J. (1972). *Semantic theory.* New York: Harper & Row.
Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.
Kosslyn, S.M., Cave, C.B., Provost, D.A., & von Gierke, S.M. (1988). Sequential processes in image generation. *Cognitive Psychology, 20,* 319-343.
Kripke, S. (1977). Speaker reference and semantic reference. In P.A. French, T.E. Uehling, Jr., & Wettstein, H.K. (Eds.), *Contemporary perspectives in the philosophy of language* (6-27). Minneapolis: University of Minnesota Press.
Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago: University of Chicago Press.
Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago: University of Chicago Press.
Langacker, R.W. (1986). An introduction to cognitive grammar. *Cognitive Science, 10,* 1-40.
Langacker, R.W. (1987). *Foundations of cognitive grammar: Vol. 1. Theoretical prerequisites.* Stanford, CA: Stanford University Press.

Malt, B.C. (1991). Word meaning and word use. In P.J. Schwanenflugel (Ed.), *The psychology of word meanings,* (37-70). Hillsdale, NJ: Erlbaum.
Mandler, J.M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review, 99,* 587-604.
Mandler, J.M., & Ritchey, G.H. (1977). Long-term memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 386-396.
Markman, A.B., & Gentner, D. (in press-a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language.*
Markman, A.B., & Gentner, D. (in press-b). Structural alignments during similarity comparisons. *Cognitive Psychology.*
McNeill, D. (1992). *Hand and mind: What gestures reveal about thought.* Chicago: University of Chicago Press.
Medin, D.L., & Ross, B.H. (1989). The specific character of abstract thought: Categorization, problem solving, and induction. In R.J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 5, 189-223). Hillsdale, NJ: Erlbaum.
Medin, D.L., & Schaffer, M. (1978). A context theory of classification learning. *Psychological Review, 85,* 207-238.
Minsky, M.L. (1977). A framework for representing knowledge. In P.H. Winston (Ed.), *The psychology of computer vision* (211-277). New York: McGraw-Hill.
Murphy, G.L. (1988). Comprehending complex concepts. *Cognitive Science, 12,* 529-562.
Murphy, G.L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language, 29,* 259-288.
Murphy, G.M. (1991). Meaning and concepts. In P.J. Schwanenflugel (Ed.), *The psychology of word meanings,* (11-35). Hillsdale, NJ: Erlbaum.
Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289-316.
Murphy, G.L., & Wisniewski, E.J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 572-586.
Nickerson, R.S., & Adams, M.J. (1979). Long-term memory for a common object. *Cognitive Psychology, 11,* 287-307.
Newtson, D. (1976). Foundations of attribution: The perception of ongoing behavior. In J.H. Harvey, W.J. Ickes, & R.F. Kidd (Eds.), *New directions in attribution research* (Vol. 1, 223-247). Hillsdale, NJ: Erlbaum.
Nosofsky, R.M. (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 54-65.
Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy, 3,* 143-184.
Palmer, S.E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition, 3,* 519-526.
Potter, M.C., & Faulconer, B.A. (1975). Time to understand pictures and words. *Nature, 253,* 437-438.
Potter, M.C., Valian, V.V., & Faulconer, B.A. (1977). Representation of a sentence and its pragmatic implications: Verbal, imagistic, or abstract? *Journal of Verbal Learning and Verbal Behavior, 16,* 1-12.
Potter, M.C., Kroll, J.F., Yachzel, B., Carpenter, E., & Sherman, J. (1986). Pictures in sentences: Understanding without words. *Journal of Experimental Psychology: General, 115,* 281-294.
Putnam, H. (1973). Meaning and reference. *Journal of Philosophy, 70,* 699-711.
Rosch, E., & Lloyd, B. (Eds.) (1978). *Cognition and categorization.* Hillsdale, NJ: Erlbaum.
Rosch, E., & Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7,* 573-605.
Ross, B.H., Perkins, S.J., & Tenpenny, P.L (1990). Reminding-based category learning. *Cognitive Psychology, 22,* 460-492.
Ross, B.H., & Spalding, T.L. (1991). Some influences of instance comparisons in concept formation. In D. Fisher & M. Pazzani (Eds.), *Computational approaches to concept formation.* San Mateo, CA: Morgan Kaufmann.
Roth, J.D., & Kosslyn, S.M. (1988). Construction of the third dimension in mental imagery. *Cognitive Psychology, 20,* 344-361.
Roth, E.M., & Shoben, E.J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15,* 346-378.

Rumelhart, D.E., & McClelland, J.L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review, 89,* 60-94.
Sachs, J.D.S. (1974). Memory in reading and listening to discourse. *Memory & Cognition, 2,* 95-100.
Schank, R.C. (1975). *Conceptual information processing.* Amsterdam: North-Holland.
Schank, R.C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people.* New York: Cambridge University Press.
Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Erlbaum.
Schubert, L.K., & Pelletier, F.J. (1987). Problems in the representation of the logical form of generics, plurals, and mass nouns. *New directions in semantics* (385-451). London: Academic Press.
Searle, J.R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3,* 417-424.
Shepard, R.N., & Cooper, L.A. (1982). *Mental images and their transformations.* New York: Cambridge University Press.
Smith, E.E. (1978). Theories of semantic memory. In W.K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 6). Hillsdale, NJ: Erlbaum.
Smith, E.E., & Medin, D.L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.
Smith, E.E., Osherson, D.N., Rips, L.J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science, 12,* 485-528.
Stein, N.L., & Levine, L.J. (1990). Making sense out of emotion: The representation and use of goal-structured knowledge. In N.L. Stein, B. Leventhal, T. Trabasso (Eds.), *Psychological and biological approaches to emotion* (45-73). Hillsdale, NJ: Erlbaum.
Stein, N.L., Trabasso, T., & Liwag, M. (in press). The representation and organization of emotional experience: Unfolding the emotional episode. In M. Lewis & J. Haviland (Eds.), *Handbook of emotion.* New York: Guilford.
Sweetser, E. (1990). *From etymology to pragmatics: Metaphorical and cultural aspects of semantics*. Cambridge, England: Cambridge University Press.
Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science, 12,* 49-100.
Thorndyke, P.W., & Hayes-Roth, B. (1979). The use of schemata in the acquisition and transfer of knowledge. *Cognitive Psychology, 11,* 82-106.
Trabasso, T., & Sperry, L.L. (1985). The causal basis for deciding importance of story events. *Journal of Memory and Language, 24,* 595-611.
Trabasso, T., & van den Broek, P. (1985). Causal thinking and story comprehension. *Journal of Memory and Language, 24,* 612-630.
Trabasso, T., van den Broek, P., & Suh, S.Y. (1989). Logical necessity and transitivity of causal relations in stories. *Discourse Processes, 12,* 1-25.
Tversky, B. (1991). Spatial mental models. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, 109-145). New York: Academic Press.
van Dijk, T.A., & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic.
Watkins, M.J., & Kerkar, S.P. (1985). Recall of a twice-presented item without recall of either presentation: Generic memory for events. *Journal of Memory and Language, 24,* 666-678.
Wellman, H.M., & Gelman, S.A. (1988). Children's understanding of the non-obvious. In R. Sternberg (Ed.), *Advances in the psychology of human intelligence* (99-135). Hillsdale, NJ: Erlbaum.
Wittgenstein, L. (1953). *Philosophical investigations* (G.E.M. Anscombe, Trans.). New York: Macmillan.