

Concepts and Categories: Memory, Meaning, and Metaphysics

Douglas L. Medin and Lance J. Rips

Introduction

The concept of concepts is difficult to define, but no one doubts that concepts are fundamental to mental life and human communication. Cognitive scientists generally agree that a concept is a mental representation that picks out a set of entities, or a category. That is, concepts *refer*, and what they refer to are categories. It is also commonly assumed that category membership is not arbitrary but rather a principled matter. What goes into a category belongs there by virtue of some law-like regularities. But beyond these sparse facts, the concept CONCEPT is up for grabs. As an example, suppose you have the concept TRIANGLE represented as “a closed geometric form having three sides.” In this case, the concept is a definition. But it is unclear what else might be in your triangle concept. Does it include the fact that geometry books discuss them (though some don’t) or that they have 180 degrees (though in hyperbolic geometry none do)? It is also unclear how many concepts have definitions or what substitutes for definitions in ones that don’t.

Our goal in this chapter is to provide an overview of work on concepts and categories in the last half century. There has been such a consistent stream of research over this period that one reviewer of this literature, Gregory Murphy (2002), felt compelled to call his monograph, *The Big Book of Concepts*. Our task is eased by recent reviews, including Murphy’s aptly named one (e.g., Medin, Lynch & Solomon, 2000; Murphy, 2002; Rips, 2001; Wisniewski, 2002). Their thoroughness gives us the luxury of doing a review focused on a single perspective or “flavor” — the relation between concepts, memory, and meaning.

Concepts and Categorization / 2

The remainder of this chapter is organized as follows. In the rest of this section, we briefly describe some of the tasks or functions that cognitive scientists have expected concepts to perform. This will provide a roadmap to important lines of research on concepts and categories. Next, we return to developments in the late 1960's and early 1970's that raised the exciting possibility that laboratory studies could provide deep insights into both concept representations and the organization of (semantic) memory. Then we describe the sudden collapse of this optimism and the ensuing lines of research that, however intriguing and important, essentially ignored questions about semantic memory. Next we trace a number of relatively recent developments under the somewhat whimsical heading, "Psychometaphysics." This is the view that concepts are embedded in (perhaps domain-specific) theories. This will set the stage for returning to the question of whether research on concepts and categories is relevant to semantics and memory organization. We'll use that question to speculate about future developments in the field. In this review, we'll follow the usual conventions of using words in all caps to refer to concepts and quoted words to refer to linguistic expressions.

Functions of concepts. For purposes of this review, we will collapse the many ways people can use concepts into two broad functions: categorization and communication. The conceptual function that most research has targeted is *categorization*, the process by which mental representations (concepts) determine whether or not some entity is a member of a category. Categorization enables a wide variety of subordinate functions because classifying something as a category member allows people to bring their knowledge of the category to bear on the new instance. Once people categorize some novel entity, for example, they can use relevant knowledge for *understanding* and *prediction*. Recognizing a cylindrical object as a flashlight allows you to understand its parts, trace its functions, and predict its behavior. For example, you can confidently infer that the flashlight will have one or more batteries, will have some sort of switch, and will normally produce a beam of light when the switch is pressed.

Concepts and Categorization / 3

Not only do people categorize in order to understand new entities, they also use the new entities to modify and update their concepts. In other words, categorization supports *learning*. Encountering a member of a category with a novel property—for example, a flashlight that has a siren for emergencies—can result in that novel property being incorporated into the conceptual representation. In other cases, relations between categories may support inference and learning. For example, finding out that flashlights can contain sirens may lead you to entertain the idea that cell phones and fire extinguishers might also contain sirens. Hierarchical conceptual relations support both inductive and deductive reasoning. If all trees contain xylem and hawthorns are trees, then one can deduce that hawthorns contain xylem. In addition, finding out that white oaks contain phloem provides some support for the inductive inference that other kinds of oaks contain phloem. People also use categories to *instantiate goals in planning* (Barsalou, 1983). For example, a person planning to do some night fishing might create an ad hoc concept, THINGS TO BRING ON A NIGHT FISHING TRIP, which would include a fishing rod, tackle box, mosquito repellent, and a flashlight.

Concepts are also centrally involved in *communication*. Many of our concepts correspond to lexical entries, such as the English word “flashlight.” In order for people to avoid misunderstanding each other, they must have comparable concepts in mind. If A’s concept of cell phone corresponds with B’s concept of flashlight, it won’t go well if A asks B to make a call. An important part of the function of concepts in communication is their ability to combine in order to create an unlimited number of new concepts. Nearly every sentence you encounter is new—one you’ve never heard or read before—and concepts (along with the sentence’s grammar) must support your ability to understand it. Concepts are also responsible for more ad hoc uses of language. For example, from the base concepts of TROUT and FLASHLIGHT, you might create a new concept, TROUT FLASHLIGHT, which in the context of our current discussion would presumably be a flashlight used when trying to catch trout (and not a flashlight with a picture of a trout on it, though this may be the correct interpretation in some other context). A

Concepts and Categorization / 4

major research challenge is to understand the principles of *conceptual combination* and how they relate to communicative contexts (see Fodor, 1994, 1998; Gleitman & Papafragou, chap. 24 of this volume; Hampton, 1997; Partee, 1995; Rips, 1995; Wisniewski, 1997).

Overview. So far, we've introduced two roles for concepts: categorization (broadly construed) and communication. These functions and associated subfunctions are important to bear in mind because studying any one in isolation can lead to misleading conclusions about conceptual structure (see Solomon, Medin, & Lynch, 1999, for a review bearing on this point). At this juncture, however, we need to introduce one more plot element into the story we are telling. Presumably everything we have been talking about has implications for human memory and memory organization. After all, concepts are mental representations, and people must store these representations somewhere in memory. However, the relation between concepts and memory may be more intimate. A key part our story is what we call "the semantic memory marriage," the idea that memory organization corresponds to meaningful relations between concepts. Mental pathways that lead from one concept to another—for example, from ELBOW to ARM—represent relations like IS A PART OF that link the same concepts. Moreover, these memory relations may supply the concepts with all or part of their meaning. By studying how people use concepts in categorizing and reasoning, researchers could simultaneously explore memory structure and the structure of the mental lexicon. In other words, the idea was to unify categorization, communication (in its semantic aspects), and memory organization. As we'll see, this marriage was somewhat troubled, and there are many rumors about its break up. But we are getting ahead of our story. The next section begins with the initial romance.

A Mini-history

Research on concepts in the middle of the last century reflected a gradual easing away from behaviorist and associative learning traditions. The focus, however, remained on learning. Most of this

Concepts and Categorization / 5

research was conducted in laboratories using artificial categories (a sample category might be any geometric figure that is both red and striped) and directed at one of two questions: (a) Are concepts learned by gradual increases in associative strength, or is learning all-or-none (Levine, 1962; Trabasso & Bower, 1968)? and (b) Which kinds of rules or concepts (e.g., disjunctive, such as RED OR STRIPED, versus conjunctive, such as RED AND STRIPED) are easiest to learn (Bruner, Goodnow, & Austin, 1956; Bourne, 1970; Restle, 1962)?

This early work tended either to ignore real world concepts (Bruner et al. represent something of an exception here) or to assume implicitly that real world concepts are structured according to the same kinds of arbitrary rules that defined the artificial ones. According to this tradition, category learning is equivalent to finding out the definitions that determine category membership.

Early Theories of Semantic Memory

Although the work on rule learning set the stage for what was to follow, two developments associated with the emergence of cognitive psychology dramatically changed how people thought about concepts.

Turning point 1: Models of memory organization. The idea of programming computers to do intelligent things (artificial intelligence or AI) had an important influence on the development of new approaches to concepts. Quillian (1967) proposed a hierarchical model for storing semantic information in a computer that was quickly evaluated as a candidate model for the structure of human memory (Collins & Quillian, 1969). Figure 1 provides an illustration of part of a memory hierarchy that is similar to what the Quillian model suggests.

Insert Figure 1 about here

First, note that the network follows a principle of cognitive economy. Properties true of all animals, like eating and breathing, are stored only with the animal concept. Similarly, properties that are generally true of birds are stored at the bird node, but properties distinctive to individual kinds (e.g., being yellow) are stored with the specific concept nodes they characterize (e.g., CANARY). A property does not have to be true of all subordinate concepts to be stored with a superordinate. This is illustrated in Figure 1, where CAN FLY is associated with the bird node; the few exceptions (e.g., flightlessness for ostriches) are stored with particular birds that do not fly. Second, note that category membership is defined in terms of positions in the hierarchical network. For example, the node for CANARY does not directly store the information that canaries are animals; instead, membership would be “computed” by moving from the canary node up to the bird node and then from the bird node to the animal node. It is as if a deductive argument is being constructed of the form, “All canaries are birds and all birds are animals and therefore all canaries are animals.”

Although these assumptions about cognitive economy and traversing a hierarchical structure may seem speculative, they yield a number of testable predictions. Assuming that traversal takes time, one would predict that the time needed for people to verify properties of concepts should increase with the network distance between the concept and the property. For example, people should be faster to verify that a canary is yellow than to verify that a canary has feathers and faster to determine that a canary can fly than that a canary has skin. Collins and Quillian found general support for these predictions.

Turning point 2: Natural concepts and family resemblance. The work on rule learning suggested that children (and adults) might learn concepts by trying out hypotheses until they hit on the correct definition. In the early 1970’s, however, Eleanor Rosch and her associates (e.g., Rosch, 1973; Rosch &

Concepts and Categorization / 7

Mervis, 1975) argued that most everyday concepts are not organized in terms of the sorts of necessary and sufficient features that would form a (conjunctive) definition for a category. Instead, such concepts depend on properties that are generally true but need not hold for every member. Rosch's proposal was that concepts have a "family resemblance" structure: What determines category membership is whether an example has enough characteristic properties (is enough like other members) to belong to the category.

One key idea associated with this view is that not all category members are equally "good" examples of a concept. If membership is based on characteristic properties and some members have more of these properties than others, then the ones with more characteristic properties should better exemplify the category. For example, canaries but not penguins have the characteristic bird properties of flying, singing, and building a nest; so one would predict that canaries would be more typical birds than penguins. Rosch and Mervis (1975) found that people do rate some examples of a category to be more typical than others and that these judgments are highly correlated with the number of characteristic features an example possesses. They also created artificial categories conforming to family resemblance structures and produced typicality effects on learning and on goodness-of-example judgments.

Rosch and her associates (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) also argued that the family resemblance view had important implications for understanding concept hierarchies. Specifically, they suggested that the correlational structure of features (instances that share some features tend to share others) created natural "chunks" or clusters of instances that correspond to what they referred to as *basic level categories*. For example, having feathers tends to correlate with nesting in trees (among other features) in the animal kingdom, and having gills with living in water. The first cluster tends to isolate birds, while the second picks out fish. The general idea is that these basic level categories provide the best compromise between maximizing within-category similarity (birds tend to be quite similar to each other) and minimizing between-category similarity (birds tend to be dissimilar to fish). Rosch et al. showed that basic level categories are preferred by adults in naming objects, are learned first

by children, are associated with the fastest categorization reaction times, and have a number of other properties that indicate their special conceptual status.

Turning Points 1 and 2 are not unrelated. To be sure, the Collins and Quillian model, as initially presented, would not predict typicality effects (but see Collins & Loftus, 1975), and it wasn't obvious that it contained anything that would predict the importance of basic level categories. Nonetheless these conceptual breakthroughs led to an enormous amount of research premised on the notion that concepts are linked in memory by meaningful pathways, so that memory groups concepts according to their similarity in meaning (see Anderson & Bower, 1973; and Norman & Rumelhart, 1975, for theories and research in this tradition, and Goldstone & Son, chap. 1 of this volume, for current theories of similarity).

Fragmentation of Semantics and Memory

Prior to about 1980, most researchers in this field saw themselves as investigating “semantic memory”—the way that long-term memory organizes meaningful information. Around 1980, the term itself became passé, at least for this same group of researchers, and the field regrouped under the banner of “Categories and Concepts” (the title of Smith & Medin's, 1981, synthesis of research in this area). At the time, these researchers may well have seen this change as a purely nominal one, but we suspect it reflected a retreat from the claim that semantic-memory research had much to say about either semantics or memory. How did this change come about?

Memory organization. Initial support for a Quillian-type memory organization came from Quillian's own collaboration with Allan Collins (Collins & Quillian, 1969), which we mentioned earlier. Related evidence also came from experiments on lexical priming: Retrieving the meaning of a word made it easier to retrieve the meaning of semantically related words (e.g., Meyer & Schvanevelt, 1971). In these lexical decision tasks, participants viewed a single string of letters on each trial and decided, under reaction time instructions, whether the string was a word (e.g., “daisy”) or a nonword (“raiszy”). The key result was that participants were faster to identify a string as a word if it followed a semantically

related item than an unrelated one. For example, reaction time for “daisy” was faster if on the preceding trial the participant had seen “tulip” than if he or she had seen “steel.” This priming effect is consistent with the hypothesis that activation from one concept spreads through memory to semantically related ones.

Later findings suggested, however, that the relation between word meaning and memory organization was less straightforward. For example, the typicality findings (see *Turning Point 2*) suggested that time to verify sentences of the form *An X is a Y* (e.g., “A finch is a bird”) might be a function of the overlap in the information that participants knew about the meaning of *X* and *Y*, rather than the length of the pathway between these concepts. The greater the information overlap—for example, the greater the number of properties that the referents of *X* and *Y* shared—the faster the time to confirm a true sentence and the slower the time to disconfirm a false one. For example, if you know a lot of common information about finches and birds but only a little common information about ostriches and birds, you should be faster to confirm the sentence “A finch is a bird” than “An ostrich is a bird.” Investigators proposed several theories along these lines that made minimal commitments to the way memory organized its mental concepts (McCloskey & Glucksberg, 1979; Smith, Shoben, & Rips, 1974; Tversky, 1977). Rosch’s (1978) theory likewise studiously avoided a stand on memory structure.

Evidence from priming in lexical decision tasks also appeared ambiguous. Although priming occurs between associatively related words (e.g., “bread” and “butter”), it is not so clear that there is priming between semantically-linked words in the absence of such associations. It’s controversial whether, for example, there is any automatic activation between “glove” and “hat,” despite their joint membership in the clothing category (see Balota, 1994, for a discussion). If memory is organized on a specifically semantic basis—on the basis of word meanings—then there should be activation between semantically related words even in the absence of other sorts of associations. A meta-analysis by Lucas

(2000) turned up a small effect of this type, but as Lucas notes, it is difficult to tell whether the semantically related pairs in these experiments are truly free of associations.

The idea that memory organization mimics semantic organization is an attractive one, and memory researchers attempted to modify the original Quillian approach to bring it into line with the results we have just reviewed (e.g., Collins & Loftus, 1975). The data from the sentence verification and lexical decision experiments, however, raised doubts about these theories. Later in this chapter we'll consider whether newer techniques can give us a better handle on the structure of memory, but for now let's turn to the other half of the memory = meaning equation.

Semantics. Specifying the meaning of individual words is one of the goals of semantics, but only one. Semantics must also account for the meaning of phrases, sentences, and longer units of language. One problem in using a theory like Quillian's as a semantic theory is how to extend its core idea—that the meaning of a word is the coordinates of a node in memory structure—to explain how people understand meaningful phrases and sentences. Of course, Quillian's theory and its successors can tell us how we understand sentences that correspond to pre-existing memory pathways. We've already seen how the model can explain our ability to confirm sentences like "A daisy is a flower." But what about sentences that don't correspond to pre-existing connections, sentences like "Fred placed a daisy in a lunchbox"?

The standard approach to sentence meaning in linguistics is to think of the meaning of sentences as built from the meaning of the words that compose them, guided by the sentence's grammar (e.g., Chierchia & McConnell-Ginet, 1990). We can understand sentences that we have never heard or read before, and since there are an enormous number of such novel sentences, we can't learn their meaning as single chunks. It therefore seems quite likely that we compute the meaning of these new sentences. But if word meaning is the position of a node in a network, it is hard to see how this position could combine with other positions to produce sentence meanings. What is the process that could take the relative

network positions for FRED, PLACE, DAISY, IN, and LUNCHBOX and turn them into a meaning for “Fred placed a daisy in a lunchbox”?

If you like the notion of word meaning as relative position, then one possible solution to the problem of sentence meaning is to connect these positions with further pathways. Since we already have an array of memory nodes and pathways at our disposal, why not add a few more in order to encode the meaning of a new sentence? Perhaps the meaning of “Fred placed a daisy in the lunchbox” is given by a new set of pathways that interconnect the nodes for FRED, PLACE, DAISY, and so on, in a configuration corresponding to the sentence’s structure. This is the route that Quillian and his successors took (e.g., Anderson & Bower, 1973; Norman & Rumelhart, 1975; Quillian, 1969), but it comes at a high price. Adding new connections changes the overall network configuration and thereby alters the meaning of the constituent terms. (Remember: meaning is supposed to be *relative* position.) But it is far from obvious that encoding incidental facts alters word meaning. It seems unlikely, for example, that learning the sentence about Fred changes the meaning of “daisy.” Moreover, because meaning is a function of the entire network, the same incidental sentences change the meaning of all words. Learning about Fred’s daisy-placing shifts the meaning of seemingly unrelated words like “hippopotamus” if only a bit.

Related questions apply to other psychological theories of meaning in the semantic-memory tradition. In order to handle the typicality results mentioned earlier, some investigators proposed that the mental representation of a category like daisies consists of a prototype for that category—for example, a description of a good example of a daisy (e.g., Hampton, 1979; McCloskey & Glucksberg, 1979). The meaning of “daisy” in these prototype theories would thus include default characteristics, such as growing in gardens, that apply to most, but not all, daisies. We will discuss prototype theories in more detail soon, but the point for now is that prototype representations for individual words are difficult to combine to obtain a meaning for phrases that contain them. One potential way to combine prototypes—fuzzy set theory (Zadeh, 1965)—proved vulnerable to a range of counterexamples (Osherson & Smith, 1981,

1982). In general, the prototypes of constituent concepts can differ from the prototypes of their combinations in unpredictable ways (Fodor, 1994). The prototype of BIRDS THAT ARE PETS (perhaps a parakeet-like bird) may differ from the prototypes of both BIRDS and PETS (see Storms, de Boeck, van Mechelen, & Ruts, 1998, for related evidence). So if word meanings are prototypes, it is hard to see how the meaning of phrases could be a compositional function of the meaning of their parts.

Other early theories proposed that category representations consist of descriptions of exemplars of the category in question. For example, the mental representation of DAISY would include descriptions of specific daisies that an individual had encoded (e.g., Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1986). However, these theories have semantic difficulties of their own (see Rips, 1995). For example, if, by chance, the only Nebraskans you've met are chiropractors and the only chiropractors you've met are Nebraskans, then exemplar models appear to mispredict that "Nebraskan" and "chiropractor" will be synonyms for you.

To recap briefly, we've found that experimental research on concepts and categories was largely unable to confirm that global memory organization (as in Quillian's semantic memory) conferred word meaning. In addition, neither the global theories that initiated this research nor the local prototype or exemplar theories that this research produced were able to provide insight into the basic semantic problem of how we understand the meaning of novel sentences. This left semantic memory theory in the unenviable position of being unable to explain either semantics or memory.

Functions and Findings

Current research in this field still focuses on categorization and communication, but without the benefit of a framework that gives a unified explanation for the functions that concepts play in categorizing, reasoning, learning, language understanding, and memory organization. In this section, we

survey the state of the art, and in the following one, we consider the possibility of reuniting some of these roles.

Category Learning and Inference

One nice aspect of Rosch and Mervis's (1975) studies of typicality effects is that they used both natural language categories and artificially-created categories. Finding typicality effects with natural (real world) categories shows that the phenomenon is of broad interest; finding these same effects with artificial categories provides systematic control for potentially confounding variables (e.g., exemplar frequency) in a way that cannot be done for lexical concepts. This general strategy linking the natural to the artificial has often been followed over the past few decades. Although researchers using artificial categories have sometimes been guilty of treating these categories as ends in themselves, there are enough parallels between results with artificial and natural categories that each area of research informs the other (see Medin & Coley, 1998 for a review).

Prototype versus exemplar models. One idea compatible with Rosch's family-resemblance hypothesis is the *prototype view*. It proposes that people learn the characteristic features (or central tendency) of categories and use them to represent the category (e.g., Reed, 1972). This abstract prototype need not correspond to any experienced example. According to this theory, categorization depends on similarity to the prototypes. For example, to decide whether some animal is a bird or a mammal, a person would compare the (representation of) that animal to both the bird and the mammal prototypes and assign it to the category whose prototype it most resembled. The prototype view accounts for typicality effects in a straightforward manner. Good examples have many characteristic properties of their category and have few characteristics in common with the prototypes of contrasting categories.

Early research appeared to provide striking confirmation of the idea of prototype abstraction. Using random dot patterns as the prototypes, Posner and Keele (1968, 1970) produced a category from each prototype. The instances in a category were "distortions" of the prototype, generated by moving

constituent dots varying distances from their original positions. Posner and Keele first trained participants to classify examples that they created by distorting the prototypes. Then they gave a transfer test in which they presented both the old patterns and new low or high distortions that had not appeared during training. In addition, the prototypes, which the participants had never seen, were presented during transfer. Participants had to categorize these transfer patterns; but unlike the training procedure, the transfer test gave participants no feedback about the correctness of their responses. The tests either immediately followed training or appeared after a one-week delay.

Posner and Keele (1970) found that correct classification of the new patterns decreased as distortion (distance from a category prototype) increased. This is the standard typicality effect. The most striking result was that a delay differentially affected categorization of prototypic versus old training patterns. Specifically, correct categorization of old patterns decreased over time to a reliably greater extent than performance on prototypes. In the immediate test, participants classified old patterns more accurately than prototypes; but in the delayed test, accuracy on old patterns and prototypes was about the same. This differential forgetting is compatible with the idea that training leaves participants with representations of both training examples and abstracted prototypes, but that memory for examples fades more rapidly than memory for prototypes. The Posner and Keele results were quickly replicated by others and constituted fairly compelling evidence for the prototype view.

But this proved to be the beginning of the story rather than the end. Other researchers (e.g., Brooks, 1978; Medin & Schaeffer, 1978) put forth an *exemplar view* of categorization. Their idea was that memory for old exemplars by itself could account for transfer patterns without the need for positing memory for prototypes. On this view, new examples are classified by assessing their similarity to stored examples and assigning the new example to the category that has the most similar examples. For instance, some unfamiliar bird (e.g., a heron) might be correctly categorized as a bird not because it is similar to a bird prototype, but rather because it is similar to flamingos, storks, and other shore birds.

In general, similarity to prototypes and similarity to stored examples will tend to be highly correlated (Estes, 1986). Nonetheless, for some category structures and for some specific exemplar and prototype models, it is possible to develop differential predictions. Medin and Schaffer (1978), for example, pitted number of typical features against high similarity to particular training examples and found that categorization was more strongly influenced by the latter. A prototype model would make the opposite prediction.

Another contrast between exemplar and prototype models revolves around sensitivity to within-category correlations (Medin, Altom, Edelson, & Freko, 1982). A prototype representation captures what is on average true of a category but is insensitive to within-category feature distributions. For example, a bird prototype could not represent the impression that small birds are more likely to sing than large birds (unless one had separate prototypes for large and small birds). Medin et al. (1982) found that people are sensitive to within category correlations (see also Malt & Smith, 1984, for corresponding results with natural object categories). Exemplar theorists were also able to show that exemplar models could readily predict other effects that originally appeared to support prototype theories—differential forgetting of prototypes versus training examples, and prototypes being categorized as accurately or more accurately than training examples. In short, early skirmishes strongly favored exemplar models over prototype models. Parsimony suggested no need to posit prototypes if stored instances could do the job. Since the early 1980's, there have been a number of trends and developments in research and theory with artificially constructed categories, and we will be able to give only the briefest of summaries here.

New models. There are now more contending models for categorizing artificial stimuli, and the early models have been extensively elaborated. For example, researchers have generalized the original Medin and Schaffer (1978) exemplar model to handle continuous dimensions (Nosofsky, 1986), to address the time course of categorization (Lamberts, 1995; Nosofsky & Palmeri, 1997a; Palmeri, 1997),

to generate probability estimates in inference tasks (Juslin & Persson, 2002), and to embed it in a neural network (Kruschke, 1992),

Three new kinds of classification theories have been added to the discussion: rational approaches, decision bound models, and neural network models. Anderson (1990, 1991) proposed that an effective approach to modeling cognition in general and categorization in particular is to analyze the information available to a person in the situation of interest and then to determine abstractly what an effective, if not optimal, strategy might be. This approach has led to some new sorts of experimental evidence (e.g., Anderson & Fincham, 1996; Clapper & Bower, 2002) and pointed researchers more in the direction of the inference function of categories. Interestingly, the Medin and Schaffer exemplar model corresponds to a special case of the rational model, and Nosofsky (1991) has discussed the issue of whether the rational model adds significant explanatory power. However, there is also some evidence undermining the rational model's predictions concerning inference (e.g., Murphy & Ross, 1994; Malt, Ross, & Murphy, 1995; Ross & Murphy, 1996; Palmeri, 1999).

Decision bound models (e.g. Ashby & Maddox, 1993; Maddox & Ashby, 1993) draw their inspiration from psychophysics and signal detection theory. Their primary claim is that category learning consists of developing decision bounds around the category that will allow people to categorize examples successfully. The closer an item is to the decision bound the harder it should be to categorize. This framework offers a new perspective on categorization in that it may lead investigators to ask questions such as: How do the decision bounds that humans adopt compare with what is optimal? What kinds of decision functions are easy or hard to acquire? Researchers have also directed efforts to distinguish decision-bound and exemplar models (e.g., McKinley & Nosofsky, 1995; Maddox & Ashby, 1998; Maddox, 1999; Nosofsky & Palmeri, 1997b; Nosofsky, 1998). One possible difficulty with decision bound models is that they contain no obvious mechanism by which stimulus familiarity can affect performance, contrary to empirical evidence that it does (Verguts, Storms, & Tuerlinckx, 2001).

Neural network or connectionist models are the third type of new model on the scene (see Knapp & Anderson, 1984, and Kruschke, 1992, for examples, and Doumas & Hummel, chap. 20 in this volume, for further discussion of connectionism). It may be a mistake to think of connectionist models as comprising a single category, as they take many forms depending on assumptions about hidden units, attentional processes, recurrence, and the like. There is one sense in which neural network models with hidden units may represent a clear advance on prototype models: They can form prototypes in a bottom-up manner that reflects within-category structure (e.g., Love, Medin, & Gureckis, in press). That is, if a category comprises two distinct clusters of examples, network models can create a separate hidden unit for each chunk (e.g., large birds versus small birds) and thereby show sensitivity to within-category correlations.

Mixed models and multiple categorization systems. A common response to hearing about various models of categorization is to suggest that all the models may be capturing important aspects of categorization and that research should determine in which contexts one strategy versus another is likely to dominate. One challenge to this divide and conquer program is that the predictions of alternative models tend to be highly correlated and separating them is far from trivial. Nonetheless, there is both empirical research (e.g. Reagher & Brooks, 1993; Nosofsky, Clark, & Shin, 1989; Johansen & Palmeri, in press) and theoretical modeling that support the idea that mixed models of categorization are useful and perhaps necessary. Current efforts combine rules and examples (e.g., Nosofsky, Palmeri, & McKinley, 1994; Erickson & Kruschke, 1998), as well as rules and decision bounds (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Some models also combine exemplars and prototypes (e.g. Homa, Sterling, & Trepal, 1981; Smith, Murray, & Minda, 1997; Smith & Minda, 1998, 2000; Minda & Smith, 2001), but it remains controversial whether the addition of prototypes is needed (e.g., Busemeyer, Dewey, & Medin, 1984; Nosofsky & Johansen, 2000; Nosofsky & Zaki, 2002; Stanton, Nosofsky, & Zaki, 2002).

The upsurge of cognitive neuroscience has reinforced the interest in multiple memory systems. One intriguing line of research by Knowlton, Squire, and associates (Knowlton & Squire, 1993; Squire & Knowlton, 1995; Knowlton, Mangels & Squire, 1996), favoring multiple categorization systems, involves a dissociation between categorization and recognition. Knowlton and Squire (1993) used the Posner and Keele dot pattern stimuli to test amnesic and matched control patients on either categorization learning and transfer or on a new-old recognition task (involving five previously studied patterns versus five new patterns). The amnesiacs performed very poorly on the recognition task but were not reliably different from control participants on the categorization task. Knowlton and Squire took this as evidence for a two system model, one based on explicit memory for examples and one based on an implicit system (possibly prototype abstraction). On this view, amnesiacs have lost access to the explicit system but can perform the classification task using their intact implicit memory.

These claims have provoked a number of counterarguments. First, Nosofsky and Zaki (1998) showed that a single system (exemplar) model could account for both types of data from both groups (by assuming that the exemplar-based memory of amnesiacs was impaired but not absent). Second, investigators have raised questions about the details of Knowlton and Squire's procedures. Specifically, Palmeri and Flanery (1999) suggested that the transfer tests themselves may have provided cues concerning category membership. They showed that undergraduates who had never been exposed to training examples (the students thought they were being shown patterns subliminally) performed above chance on transfer tests in this same paradigm. The debate is far from resolved, and there are strong advocates both for and against the multiple systems view (e.g. Reber, Stark, & Squire, 1998a, b; Nosofsky & Johansen, 2000; Filoteo, Maddox, & Davis, 2002; Maddox, 2002; Palmeri & Flanery, 2002). It is safe to predict that this issue will receive continuing attention.

Inference learning. Recently investigators have begun to worry about extending the scope of category learning studies by looking at inference. Often we categorize some entity in order to help us

accomplish some function or goal. Ross (1997, 1999, 2000) has shown that the category representations people develop in laboratory studies depend on use and that use affects later categorization. In other words, models of categorization ignore inference and use at their peril. Other work suggests that having a cohesive category structure is more important for inference learning than it is for classification (Yamauchi & Markman, 1998, 2000a, 2000b; Yamauchi, Love, & Markman, 2002; for modeling implications see Love, Markman, & Yamauchi, 2000; Love, et al., in press). More generally, this work raises the possibility that diagnostic rules based on superficial features, which appear so prominently in pure categorization tasks, may not be especially relevant for contexts involving multiple functions or more meaningful stimuli (e.g., Markman & Makin, 1998; Wisniewski & Medin, 1994).

Feature learning. The final topic on our “must mention” list for work with artificial categories is feature learning. It is a common assumption in both models of object recognition and category learning that the basic units of analysis or features remain unchanged during learning. There is increasing evidence and supporting computational modeling indicating that this assumption is incorrect. Learning may increase or decrease the distinctiveness of features and may even create new features (see Goldstone, 1998, 2003; Goldstone and Steyvers, 2001; Goldstone, Lippa, & Shrifin, 2001; Schyns & Rodet, 1997; Schyns, Goldstone, & Thibaut, 1998).

Feature learning has important implications for our understanding of the role of similarity in categorization. It is intuitively compelling to think of similarity as a causal factor supporting categorization—things belong to the same category because they are similar. But this may have things backwards. Even standard models of categorization assume that learners selectively attend to features that are diagnostic, and the work on feature learning suggests that learners may create new features that help partition examples into categories. In that sense, similarity (in the sense of overlap in features) is the byproduct, not the cause, of category learning. We’ll take up this point again in discussing the theory of categorization later in this review.

Reasoning. As we noted earlier, one of the central functions of categorization is to support reasoning. Having categorized some entity as a bird, one may predict with reasonable confidence that it builds a nest, sings, and can fly, though none of these inferences is certain. In addition, between-category relations may guide reasoning. For example, from the knowledge that robins have some enzyme in their blood one is likely to be more confident that the enzyme is in sparrows than in raccoons. The basis for this confidence may be that robins are more similar to sparrows than to raccoons or that robins and sparrows share a lower rank superordinate category than do robins and raccoons (birds versus vertebrates). We will not review this literature here because Sloman and Lagnado (chapter 3) summarize it nicely.

Summary. Bowing to practicalities, we have glossed a lot of research and skipped numerous other relevant studies. The distinction between artificially-created and natural categories is itself artificial, at least in the sense that it has no clear definition or marker. When we take up the idea that concepts may be organized in terms of theories, we will return to some laboratory studies that illustrate this fuzzy boundary. For the moment, however, we shift attention to the more language-like functions of concepts.

Language Functions

Most investigators in the concepts-and-categories area continue to assume that, in addition to their role in recognition and category learning, concepts also play a role in understanding language and in thinking discursively about things. In addition to determining, for example, which perceptual patterns signal the appearance of a daisy, the DAISY concept also contributes to the meaning of sentences like our earlier example, “Fred placed a daisy in a lunchbox.” We noted that early psychological research on concepts ran into problems in explaining the meaning of linguistic units larger than single words. Most early theories posited representations, such as networks, exemplars, or prototypes, that didn’t combine easily and, thus, complicated the problem of sentence meaning. Even if we reject the idea that sentence meanings are compositional functions of word meaning, we still need a theory of sentence meanings, and

no obvious contenders are in sight. In this section, we return to the role that concepts play in language understanding to see whether new experiments and theories have clarified this relationship

Concepts as positions in memory structures. One difficulty with the older semantic-memory view of word meaning is that memory seems to change with experience from one person to another, while meaning must be more or less constant. The sentences that you have encoded about daisies may differ drastically from those that we have encoded, since your conversation, reading habits, and other verbal give-and-take can diverge in important ways from ours. If meaning depends on memory for these sentences, then your meaning for “daisy” should likewise differ from ours. This raises the question of how you could possibly understand the sentences in this chapter in the way we intend or how you could meaningfully disagree with us about some common topic (see Fodor, 1994).

It’s possible that two people—say, Calvin and Martha—might be able to maintain mutual intelligibility as long as their conceptual networks are not too different. It is partly an empirical question how much their networks can vary while still allowing Calvin’s concepts to map correctly into Martha’s. To investigate this issue, Goldstone and Rogosky (2002) have recently carried out some simulations that try to recover such a mapping. The simulations modeled Calvin’s conceptual system as the distance between each pair of his concepts (e.g., the distance between DOG and CAT in Calvin’s system might be one unit, while the distance between DOG and DAISY might be six units). Martha’s conceptual system was represented in the same way (i.e., by exactly the same interconcept distances), except for random noise that Goldstone and Rogosky added to each distance to simulate the effect of disparate beliefs. A constraint-satisfaction algorithm then applied to Calvin’s and Martha’s systems that attempted to recover the original correspondence between the concepts—to map Calvin’s DOG to Martha’s DOG, Calvin’s DAISY to Martha’s DAISY, and so on. The results of the stimulations show that with 15 concepts in each system (the maximum number considered and the case in which the model performed best) and with no noise added to Martha’s system, the algorithm was always able to find the correct correspondence.

When the simulation added to each dimension of the interconcept distance in Martha a small random increment (drawn from a normal distribution with mean 0 and standard deviation equal to .004 times the maximum distance), the algorithm recovered the correspondence about 63% of the time. When the standard deviation increased to .006 times the maximum distance, the algorithm succeeded about 15% of the time (Goldstone & Rogosky, 2002, Figure 2).

What should one make of the Goldstone and Rogosky results? Correspondences may be recovered for small amounts of noise, but performance trailed off dramatically for larger amounts of noise. Foes of the meaning-as-relative-position theory might claim that the poor performance under the .6% noise condition proves their contention. Advocates would point to the successful part of the simulations and note that its ability to detect correct correspondences usually improved as the number of points increased (although there are some nonmonotonicities in the simulation results that qualify this finding). Clearly, this is only the beginning of the empirical side of the debate. For example, the differences between Martha and Calvin are likely to be not only random, but also systematic, as in the case where Martha grew up on a farm and Calvin was a city kid.

Concept combination. Let's look at attempts to tackle head-on the problem of how word-level concepts combine to produce the meanings of larger linguistic units. There is relatively little research in this tradition on entire sentences (see Conrad & Rips, 1986; Rips, Shoben, & Smith, 1978), but there has been a fairly steady research stream devoted to noun phrases, including adjective-noun ("edible flowers"), noun-noun ("food flowers"), and noun-relative-clause combinations ("flowers that are foods"). We'll call the noun or adjective parts of these phrases *components*, and distinguish the main or *head noun* ("flowers" in each of our examples) from the adjective or noun *modifier* ("edible" or "food"). The aim of the research in question is to describe how people understand these phrases and, in particular, how the typicality of an instance in these combinations depends on the typicality of the same instance in the components. How does the typicality of a marigold in the category of edible flowers depend on the

typicality of marigolds in the categories of edible things and flowers? As we have already noticed, this relationship is far from straightforward (parakeets are superbly typical as pet birds but less typical pets and even less typical birds).

There is an optimistic way of looking at the results of this research program and a pessimistic way, as well (for recent, mostly optimistic, reviews of this work, see Hampton, 1997; Murphy, 2002; Rips, 1995; and Wisniewski, 1997). The optimistic angle is that interesting phenomena have turned up in investigating the typicality structure of combinations. The pessimistic angle, which is a direct result of the same phenomena, is that little progress has been made in figuring out a way to predict the typicality of a combination from the typicality of its components. This difficulty is instructive, in part because all psychological theories of concept combination posit complex, structured representations, and they depict concept combination either as rearranging (or augmenting) the structure of the head noun by means of the modifier (Franks, 1995; Murphy, 1988; Smith, Osherson, Rips, & Keane, 1988) or as fitting both head and modifier into a larger relational complex (Gagné & Shoben, 1997). Table 1 summarizes what's on offer from these theories. Earlier models (at the top of the table) differ from later ones mainly in terms of the complexity of the combination process. Smith et al. (1988), for example, aimed at explaining simple adjective-noun combinations (e.g., "white vegetable") that, roughly speaking, refer to the intersection of the sets denoted by modifier and head (white vegetables are approximately the intersection of white things and vegetables). In this theory, combination occurs when the modifier changes the value of an attribute in the head noun (changing the value of the color attribute in VEGETABLE to WHITE) and boosts the importance of this attribute in the overall representation. Later theories attempted to account for nonintersective combinations (e.g., "criminal lawyers," who are often not both criminals and lawyers). These combinations call for more complicated adjustments, for example, determining a relation that links the modifier and head (a criminal lawyer is a lawyer whose clients are in for criminal charges) or

extracting a value from the modifier that can then be assigned to the head (e.g., a panther lawyer might be one who is especially vicious or tenacious).

Insert Table 1 about here

So why no progress? One reason is that many of the combinations that investigators have studied are familiar or, at least, have familiar referents. Some people have experience with edible flowers, for example, and know that they include nasturtiums, are sometimes used in salads, are often brightly colored, are peppery-tasting, and so on. We learn many of these properties by direct or indirect observation (by what Hampton, 1987, calls “extensional feedback”), and they are sometimes impossible to learn simply by knowing the meaning of “edible” and “flower.” Because these properties can affect the typicality of potential instances, the typicality of these familiar combinations won’t be a function of the typicality of their components. This means that if we are going to be able to predict typicality in a compositional way, we will have to factor out the contribution of these directly acquired properties. Rips (1995) refers to this filtering as the “No Peeking Principle”—no peeking at the referents of the combination. Of course, you might be able to predict typicality if you already know the relevant real-world facts in addition to knowing the meaning of the component concepts. The issue about understanding phrases, however, is how we are able to interpret an unlimited number of new ones. For this purpose, people need some procedure for computing new meanings from old ones that is not restricted by the limited set of facts they happened to have learned (e.g., through idiosyncratic encounters with edible flowers).

Another reason for lack of progress is that some of the combinations used in this research may be compounds or lexicalized phrases (e.g., “White House” [accent on “White”] = the residence of the President) rather than modifier-head constructions (e.g., “white house” [accent on “house”] = a house

whose color is white). Compounds are often idiomatic; their meaning is not an obvious function of their parts (see Gleitman & Gleitman's, 1970, distinction between phrasal and compound constructions; and Partee, 1995).

There is a deeper reason, however, for the difficulty in predicting compound typicality from component typicality. Even if we adhere to the No Peeking Principle, and even if we stick to clear modifier-head constructions, the typicality of a combination can depend on “emergent” properties that are not part of the representation of either component (Hastie, Schroeder, & Weber, 1990; Kunda, Miller, & Claire, 1990; Medin & Shoben, 1988; Murphy, 1988). For example, you may never have encountered, or even thought about, a smoky apple (so extensional feedback does not inform your conception of the noun phrase), but nevertheless it is plausible to suppose that smoky apples are not good tasting. Having a bad taste, however, is not a usual property of (and is not likely to be stored as part of a concept for) either apples or smoky things; on the contrary, many apples and smoky things (e.g., smoked meats, cheese, and fish) are often quite good tasting. If you agree with our assessment that smoky apples are likely to be bad tasting, that's probably because you imagine a way in which apples could become smoky (being caught in a kitchen fire, perhaps) and you infer that under these circumstances the apple wouldn't be good to eat. The upshot is that the properties of a combination can depend on complex inductive or explanatory inferences (Johnson & Keil, 2000; Kunda et al., 1990). If these properties affect the typicality of an instance with respect to the combination, then there is little hope of a simple model of this phenomenon. No current theory comes close to providing an adequate and general account of these processes.

Inferential vs. atomistic concepts. Research on the typicality structure of noun phrases is of interest for what it can tell us about people's inference and problem-solving skills. But because these processes are quite complex—drawing on general knowledge and inductive reasoning to produce emergent information—we can't predict noun phrase typicality in other than a limited range of cases. For much the same reason, typicality structure does not appear very helpful in understanding how people

construct the meaning of a noun phrase while reading or listening. By themselves, emergent properties do not rule out the possibility of a model that explains how people derive the meaning of a noun phrase from the meaning of its components. Compositionality doesn't require that all aspects of the noun phrase's meaning are parts of the components' meanings. It's sufficient to find some computable function from the components to the composite that is simple enough to account for people's understanding (see Partee, 1995, for a discussion of types of composition). The trouble is that if noun phrases' meanings require theory construction and problem solving, such a process is unlikely to explain the ease and speed with which we usually understand them in ongoing speech.

Of course, we have only considered the role of schemas or prototypes in concept combination, but it's worth noting that many of the same problems with semantic composition affect other contemporary theories, such as latent semantic analysis (Landauer & Dumais, 1997), which take a global approach to meaning. Latent semantic analysis takes as input a table of the frequencies with which words appear in specific contexts. In one application, for example, the items comprise about 60,000 word types taken from 30,000 encyclopedia entries, and the table indicates the frequency with which each word appears in each entry. The analysis then applies a technique similar to factor analysis to derive an approximately 300 dimensional space in which each word appears as a point and in which words that tend to co-occur in context occupy neighboring regions in the space. Because this technique finds a best fit to a large corpus of data, it is sensitive to indirect connections between words that inform their meaning. However, the theory has no clear way to derive the meaning of novel sentences. Although latent semantic analysis could represent a sentence as the average position of its component words, this would not allow it to capture the difference between, say, *The financier dazzled the movie star* versus

The movie star dazzled the financier, which depend on sentence structure as well as word meaning. In addition, the theory uses the distance between two words in semantic space to represent the relation between them; so the theory has trouble with semantic relations that, unlike distances, are asymmetric. It's unclear, for example, how it could cope with the fact that *father* implies *parent* but *parent* doesn't imply *father*.

On one hand, on-line sentence understanding is a rapid, reliable process. On the other, the meaning of even simple adjective-noun phrases seems to require heavy inductive inferences. Perhaps we should distinguish, then, between the *interpretation* of a phrase or sentence and its *comprehension* (Burge, 1999). On this view, comprehension gives us a more or less immediate understanding of novel phrases, based primarily on the word meaning of the components and syntactic/semantic structure. Interpretation, by contrast, is a potentially unlimited process, relying on the result of comprehension plus inference and general knowledge. The comprehension/interpretation distinction may be more of a continuum than a dichotomy, but the focus on the interpretation end of the continuum means that research on concepts is difficult to apply to comprehension. As we've just noticed, it's hard, if not impossible, to compute the typicality structure of composites. So if we want something readily computable in order to account for comprehension, we have to look to something simpler than typicality structures (and the networks, prototypes, schemas, or theories that underlie them). One possibility (Fodor, 1994, 1998) is to consider a representation in which word meanings are mental units not much different from the words themselves, and whose semantic values derive from (unrepresented) causal connections to their referents.

Generic noun phrases. Even if we abandon typicality structures as accounts of comprehension, however, it doesn't follow that these structures are useless in explaining all linguistic phenomena. Recent research on two fronts seems to us to hold promise for interactions between psychological and linguistic theories. First, there are special constructions in English that, roughly speaking, describe default

characteristics of members of a category. For example, “Lions have manes” means (approximately) that having a mane is a characteristic property of lions. Bare plural noun phrases (i.e., plurals with no preceding determiners) are one way to convey such a meaning as we’ve just noticed, but indefinite singular sentences (“A lion has a mane”) and definite singular sentences (“The lion—*Panthera leo*—has a mane”) can also convey the same idea in some of their senses. These generic sentences seem to have normative content. Unlike “Most lions have manes,” generic sentences seem to hold despite the existence of numerous exceptions; “Lions have manes” seems to be true despite the fact that most lions (e.g., female and immature lions) don’t have manes (see Krifka et al., 1995, for an introduction to generic sentences). There is an obvious relation between the truth or acceptability of generic sentences and the typicality structure of categories, since the typical properties of a category are those that appear in true generic sentences. Of course, as Krifka et al. note, this may simply be substituting one puzzle (the truth conditions of generic sentences) for another (the nature of typical properties), but this may be one place where linguistic and cognitive theories might provide mutual insight. Research by Susan Gelman and her colleagues (see Gelman, 2003, for a thorough review) suggests that generic sentences are a frequent way for caregivers to convey category information to children. Four year-olds differentiate sentences with bare plurals (“Lions have manes”) from those explicitly quantified by “all” or “some” in comprehension, production, and inference tasks (Gelman, Star, & Flukes, 2002; Hollander, Gelman, & Star, 2002). It would be of interest to know, however, at what age, and in what way, children discriminate generics from accidental generalizations—for example, when they first notice the difference between “Lions have manes” and “Lions frequently have manes” or “Most lions have manes.”

Polysemy. A second place to look for linguistic-cognitive synergy is in an account of the meanings of polysemous words. Linguists (e.g., Lyons, 1977, chap. 13) traditionally distinguish homonyms like “mold,” which have multiple unrelated meanings (e.g., a form into which liquids are poured vs. a fungus) from polysemous terms like “line” that have multiple related meanings (e.g., a

geometric line vs. a fishing line vs. a line of people, etc.). What makes polysemous terms interesting to psychologists in this area is that the relations among their meanings often possess a kind of typicality structure of their own. This is the typicality of the senses of the expression rather than the typicality of the referents of the expression, and so a type of higher-level typicality phenomenon. Figure 2 illustrates such a structure for the polysemous verb “crawl,” as analyzed by Fillmore and Atkins (2000). A rectangle in the figure represents each sense or use and includes both a brief label indicating its distinctive property and an example from British corpuses. According to Fillmore and Atkins, the central meanings for *crawl* have to do with people or creatures moving close to the ground (these uses appear in rectangles with darker outlines in the figure). But there are many peripheral uses, for example, time moving slowly (“The hours seemed to crawl by”) and creatures teeming about (“The picnic supplies crawled with ants”). The central meanings are presumably the original ones, with the peripheral meanings derived from these by a historical chaining process. Malt, Sloman, Gennari, Shi, and Wang (1999) have observed similar instances of chaining in people’s naming of artifacts, such as bottles and bowls, and it is possible that the gerrymandered naming patterns reflect the polysemy of the terms (e.g., “bottle”) rather than different uses of the same meaning. As Figure 2 shows, it’s far from easy to distinguish different related meanings (polysemy) from different uses of the same meaning (contextual variation) and from different unrelated meanings (homonymy).

Insert Figure 2 about here

Some research has attacked the issue of whether people store each of the separate senses of a polysemous term (Klein & Murphy, 2002) or store only the core meaning, deriving the remaining senses as needed for comprehension (Caramazza & Grober, 1976; Franks, 1995). Conflicting evidence in this respect may be due to the fact that some relations between senses seem relatively productive and

derivable (*regular* polysemy, such as the relationship between terms for animals and their food products, e.g., the animal meaning of “lamb” and its menu meaning), while other senses seem ad hoc (e.g., the relation between “crawl” = *moving close to the ground* and “crawl” = *teeming with people* in Figure 2). Multiple mechanisms are likely to be at work here.

Summary. We don’t mean to suggest that the only linguistic applications of psychologists’ “concepts” are in dealing with interpretation, generic phrases, and polysemy. Far from it. There are many areas, especially in developmental psycholinguistics, that hold the promise of fruitful interactions, but that we cannot review here. Nor are we suggesting that investigators in this area give up the attempt to study the use of concepts in immediate comprehension. But concepts for comprehension seem to have different properties from the concepts that figure in the other functions we have discussed, and researchers need to direct more attention to the interface between them.

Theories, Modules, and Psychometaphysics

We’ve seen, so far, some downward pressure on cognitive theories to portray human concepts as mental entities that are as simple and streamlined as possible. This pressure comes, not only from the usual goal of parsimony, but also from the role that concepts play in immediate language comprehension. But there is also a great deal of upward pressure, pressure to include general knowledge about a category as part of its representation. For example, the presence of emergent properties in concept combinations suggests that people use background knowledge in interpreting these phrases. Similarly, people may bring background knowledge and theories to bear in classifying things, even when they know a decision rule for the category. Consider psychodiagnostic classification. Although DSM-IV (the official diagnostic manual of the American Psychological Association) is atheoretical and organized in terms of rules, there is clear evidence that clinicians develop theories of disorders and, contra DSM-IV, weight

causally-central symptoms more than causally-peripheral symptoms (e.g. Kim and Ahn, in press, a). The same holds for laypeople (e.g. Furnham, 1995; Kim and Ahn, in press, b).

In this section we examine the consequences of expanding the notion of a concept to include theoretical information about a category. In the case of the natural categories, this information is likely to be causal, since people probably view physical causes as shaping and maintaining these categories. For artifacts, the relevant information may be the intentions of the person creating the object (e.g., Bloom, 1996). The issues we raise here concern the content and packaging of these causal beliefs.

The first of these issues focuses on people's beliefs about the locus of these causal forces—what we called “psychometaphysics.” At one extreme, people may believe that each natural category is associated with a single source, concentrated within a category instance, that controls the nature of that instance. The source could determine, among other things, the instance's typical properties, its category membership, and perhaps even the conditions under which it comes into and goes out of existence. Alternatively, people may believe that the relevant causal forces are more like a swarm—not necessarily internal to an instance, nor necessarily emanating from a unitary spot—but shaping the category in aggregate fashion.

The second issue has to do with the cognitive divisions that separate beliefs about different sorts of categories. People surely think that the causes that help shape daisies differ in type from those that shape teapots. Lay theories about flowers and other living things include at least crude information about specifically biological properties, whereas lay theories of teapots and other artifacts touch instead on intended and actual functions. But how deep do these divisions go? On one hand, beliefs about these domains could be modular (relatively clustered, relatively isolated), innate, universal, and local to specific brain regions. On the other, they may be free-floating, learned, culturally-specific, and distributed across cortical space. This issue is important to us because it ultimately affects whether we can patch up the “semantic memory” marriage.

Essentialism and Sortalism.

Psychological essentialism. What's the nature of people's beliefs about the causes of natural kinds? One hypothesis is that people think there is something internal to each member of the kind—an essence—that is responsible for its existence, category membership, typical properties, and other important characteristics (e.g., Atran, 1998; Gelman & Hirschfeld, 1999; Medin & Ortony, 1989). Of course, it's unlikely that people think that all categories of natural objects have a corresponding essence. There is probably no essence of pets, for example, that determines an animal's pet status. But for basic-level categories, such as dogs or gold or daisies, it's tempting to think that something in the instance determines crucial aspects of its identity. Investigators who have accepted this hypothesis are quick to point out that the theory applies to people's beliefs and not to the natural kinds themselves. Biologists and philosophers of science agree that essentialism will not account for the properties and variations that real species display, in part because the very notion of species is not coherent (e.g. Ghiselin, 1981, Hull, 1999). Chemical kinds, for example, gold, may conform much more closely to essentialist doctrine (see Sober, 1980). Nevertheless, expert opinion is no bar to laypersons' essentialist views on this topic. In addition, psychological essentialists have argued that people probably do not have a fully fleshed out explanation of what the essence is. What they have, on this hypothesis, is an IOU for a theory: a belief that there must be *something* that plays the role of essence, even though they can't supply a description of it (Medin & Ortony, 1989).

Belief in a hypothetical, minimally described essence may not seem like the sort of thing that could do important cognitive work, but psychological essentialists have pointed out a number of advantages that essences might afford, especially to children. The principle advantage may be induction potential. Medin (1989) suggested that essentialism is poor metaphysics but good epistemology in that it may lead people to expect that members of a kind will share numerous, unknown properties—an

assumption that is sometimes correct. In short, essences have a motivational role to play in getting people to investigate kinds' deeper characteristics. Essences also explain why category instances seem to run true to type—for example, why the offspring of pigs grow up to be pigs rather than cows. They also explain the normative character of kinds (e.g., their ability to support inductive arguments and their ability to withstand exceptions and superficial changes), as well as people's tendency to view terms for kinds as well-defined.

Evidence for essentialism tends to be indirect. There are results that show that children and adults do in fact hold the sorts of beliefs that essences can explain. By the time they reach first or second grade, children know that animals whose insides have been removed are no longer animals, that baby pigs raised by cows grow up to be pigs rather than cows (Gelman & Wellman, 1991), and that cosmetic surgery does not alter basic-level category membership (Keil, 1989). Research on adults also shows that “deeper” causes—those that themselves have few causes but many effects—tend to be more important in classifying than shallower causes (Ahn, 1998; Sloman, Love, & Ahn, 1998).

However, results like these are evidence for essence only if there are no better explanations for the same results, and it seems at least conceivable that children and adults make room for multiple types and sources of causes that are not yoked to an essence. According to Strevens (2000), for example, although people's reasoning and classifying suggest that causal laws govern natural kinds, it may be these laws alone, rather than a unifying essence, that are responsible for the findings. According to essentialists, people think there is something (an essence) that is directly or indirectly responsible for the typical properties of a natural kind. According to Strevens' minimalist alternative, people think that for each typical property there is something that causes it, and that something may vary for different properties. It is important to settle this difference—the presence or absence of a unique central cause—if only because the essentialist claim is the stronger one.

Essentialists counter that both children and adults assume a causal structure consistent with essence (see Kalish, 1995, 2002; Braisby, Franks, & Hampton, 1996; Diesendruck and Gelman, 1999, for debate on this issue). One strong piece of evidence for essentialism is that participants who have successfully learned artificial, family-resemblance categories (i.e., those in which category members have no single feature in common) nevertheless believe that each category contained a common, defining property (Brooks & Wood, as cited by Ahn et al., 2001). Other studies with artificial “natural” kinds have directly compared essentialist and nonessentialist structures, but have turned in mixed results (e.g., Rehder & Hastie, 2002). It’s possible that explicit training overrides people’s natural tendency to think in terms of a common cause.

In the absence of more direct evidence for essence, the essentialist-minimalist debate is likely to continue (see Ahn et al., 2001; Sloman & Malt, in press; and Strevens, 2001, for the latest salvos in this dispute). Indeed, the authors of this chapter are not in full agreement. Medin finds minimalism too unconstrained, whereas Rips opines that essentialism suffers from the opposite problem. Adding a predisposition towards parsimony to the minimalist view seems like a constructive move, but such a move would shift minimalism considerably closer to essentialism. Ultimately the issue boils down to determining to what extent causal understandings are biased toward the assumption of a unique, central cause for a category’s usual properties.

Sortalism. According to some versions of essentialism, an object’s essence determines not only which category it belongs to, but also the object’s very identity. According to this view, it’s by virtue of knowing that Fido is a dog that you know (in principle) how to identify Fido over time, how to distinguish Fido from other surrounding objects, and how to determine when Fido came into existence and when he will go out of it. In particular, if Fido happens to lose his dog essence, then Fido not only ceases to be a dog, he ceases to exist entirely. As we noted in discussing essentialism, not all categories provide these identity conditions. Being a pet, for example, doesn’t lend identify to Fido, since he may

continue to survive in the wild as a nonpet. According to one influential view (Wiggins, 1980), the critical identity-lending category is the one that answers the question *What is it?* for an object, and since basic-level categories are sometimes defined in just this way, basic-level categories are the presumed source of the principles of identity. (Theories of this type usually assume that identity conditions are associated with just one category for each object, since multiple identity conditions lead to contradictions; see Wiggins, 1980). Contemporary British philosophy tends to refer to such categories as *sortals*, however, and we will adopt this terminology here.

Sortalism plays an important role in current developmental psychology because developmentalists have used children's mastery of principles of identity to decide whether these children possess the associated concept. In some well-known studies, Xu and Carey (1996) staged for infants a scene in which (e.g.) a toy duck appears from one side of an opaque screen and then returns behind it. A toy truck next emerges from the other side of the screen and then returns to its hidden position. Infants habituate after a number of encores of this performance, at which time the screen is removed to reveal both the duck and truck (the scene that adults expect) or just one of the objects (duck or truck). Xu and Carey report that younger infants (e.g., 10 month-olds) exhibit no more surprise at seeing one object than at seeing two, while older infants (and adults) show more surprise at the one-object tableau. Xu and Carey also show in control experiments that younger and older infants perform identically if they see a preview of the two starring objects together before the start of the performance. The investigators infer that the younger infants lack the concepts DUCK and TRUCK, since they are unable to use a principle of identity for these concepts to discern that a duck cannot turn into a truck while behind the screen. Xu and Carey's experiments have sparked a controversy about whether the experimental conditions are simple enough to allow babies to demonstrate their grip on object identity (see Wilcox & Baillargeon, 1998; Xu, in press), but for present purposes what's important is the assumption that infants' inability to re-identify objects over temporal gaps implies lack of the relevant concepts.

Sortal theories impose strong constraints on some versions of essentialism. We noted that one of essentialism's strong points is its ability to explain some of the normative properties of concepts, for example, the role concepts play in inductive inferences. However, sortalism places some restrictions on this ability. Members of sortal categories can't lose their essence without losing their existence, even in counterfactual circumstances. This means that if we're faced with a premise like *Suppose dogs can bite through wire...*, we can't reason about this supposition by assuming that the essence of dogs has changed in such a way as to make dogs stronger. A dog with changed essence is not a superdog, according to sortalism, but rather has ceased to exist (see Rips, 2001). For the same reason, it is impossible to believe without contradiction both that basic-level categories are sortals and that objects can shift from one basic-level category to another.

These consequences of sortalism may be reasonable ones, but it is worth considering the possibility that sortalism—however well it fares as a metaphysical outlook—incorrectly describes people's views about object identity. Although objects typically don't survive a leap from one basic-level category to another, it may not be impossible for them to do so. Blok, Newman, and Rips (in press) and Liittschwager (1997) gave participants scenarios that described novel transformations that sometimes altered the basic level category. In both studies, participants were more likely to agree that the transformed object was identical to the original if the transformational distance was small. But these judgments could not always be predicted by basic level membership.

Results from these sci-fi scenarios should be treated cautiously, but they suggest that people think individual objects have an integrity that does not necessarily line up with their basic-level category. Although this idea may be flawed metaphysics, it is not unreasonable as psychometaphysics. People may think that individuals exist as the result of local causal forces, forces that are only loosely tethered to basic-level kinds. As long as these forces continue to support the individual's coherence, it can exist even if it finds itself in a new basic-level category. Of course, not all essentialists buy into this link between

sortalism and essentialism. For example, people might believe that individuals have both a category essence *and* a set of historical and other characteristics that makes it unique. Gutheil and Rosengren (1996) hypothesize that objects have two difference essences, one for membership and another for identity. Just how individual identity and kind identity play out under these scenarios could then be highly variable.

Domain Specificity

The notion of domain-specificity has served to organize a great deal of research on conceptual development. For example, much of the work on essentialism has been conducted in the context of exploring children's naïve biology (see also Au, 1994; Carey, 1995; Gopnik & Wellman, 1994; Spelke, Phillips, & Woodward, 1995). Learning in a given domain may be guided by certain skeletal principles, constraints, and (possibly innate) assumptions about the world (see Keil 1981; Kellman & Spelke 1983; Spelke 1990; Markman 1990; Gelman & Coley 1990; Gelman, 2003). Susan Carey's influential (1985) book presented a view of knowledge acquisition as built on framework theories that entail ontological commitments in the service of a causal understanding of real world phenomena. Two domains can be distinguished from one another if they represent ontologically distinct entities and sets of phenomena and are embedded within different causal explanatory frameworks. These ontological commitments serve to organize knowledge into domains such as naïve physics (or mechanics), naïve psychology, or naïve biology (e.g. see Wellman & Gelman 1992; Spelke et al 1995; Keil, 1994; Au, 1994; Carey 1995; Gelman & Koenig, 2001; Hatano & Inagaki, 1994; Gopnik & Wellman, 1994). In the following we will focus on one candidate domain, naïve biology.

Folkbiology and universals. There is fairly strong evidence that all cultures partition local biodiversity into taxonomies whose basic level is that of the "generic species" (Berlin et al., 1973; Atran, 1990). Generic species often correspond to scientific species (e.g., elm, wolf, and robin); however, for the large majority of perceptually salient organisms (see Hunn, 1999), such as vertebrates and flowering

plants, a scientific genus frequently has only one locally occurring species (e.g., bear, oak). In addition to the spontaneous division of local flora and fauna into generic species, cultures seem to structure biological kinds into hierarchically-organized groups, such as white oak/oak/tree. Folkbiological ranks vary little across cultures as a function of theories or belief systems (see Malt, 1994, for a review). For example, in studies with Native American and various USA and Lowland Maya groups, correlations between folk taxonomies and classical evolutionary taxonomies of the local fauna and flora average $r = .75$ at the generic-species level and about 0.5 with higher levels included (Atran, 1999; Bailenson et al., 2002; Medin et al., 2002). Much of the remaining variance owes to obvious perceptual biases (Itza' Maya group bats with birds in the same life form) and local ecological concerns. Contrary to received notions about the history and cross-cultural basis for folkbiological classification, utility does *not* appear to drive folk taxonomies (cf. Berlin et al., 1973).

These folk taxonomies also appear to guide and constrain reasoning. For example, Coley, Medin, and Atran (1997) found that both Itza' Maya and USA undergraduates privilege the generic-species level in inductive reasoning. That is, an inference from Swamp White Oak to all White Oaks is little if any stronger than an inference from Swamp White Oak to all Oaks. Above the level of Oak, however, inductive confidence takes a sharp drop. In other words, people in both cultures treat the generic level (e.g. Oak) as maximizing induction potential. The results for undergraduates are surprising, since the original Rosch et al. (1975) basic level studies had suggested that a more abstract level (e.g., TREE) acted as basic for undergraduates and should have been privileged in induction. That is, there is a discrepancy between results with undergraduates on basicness in naming, perceptual classification, and feature listing, on the one hand, and inductive inference, on the other. Coley et al. suggest that the reasoning task relies on expectations rather than knowledge and that undergraduates may know very little about biological kinds (see also Wolff, Medin, & Pankratz, 1999). Medin and Atran (in press) caution against generalizing

results on biological thought from undergraduates, since most have relatively little first-hand experience with nature.

Interdomain differences. One of the most contested domain distinctions, and one that has generated much research, is that between psychology and biology (e.g., Au & Romo, 1996, 1999; Carey 1991; Coley 1995; Gelman, in press; Hatano & Inagaki, 1996, 2001; Inagaki 1997; Inagaki & Hatano 1993, 1996; Johnson & Carey, 1998; Keil 1995; Keil et al. 1999; Rosengren et al., 1991; Springer & Keil, 1989, 1991). Carey (1985) argued that children initially understand biological concepts like ANIMAL in terms of folk psychology, treating animals as similar to people in having beliefs and desires. Others (e.g., Keil, 1989) argue that young children do have biologically-specific theories, albeit more impoverished than those of adults. For example, Springer and Keil (1989) show that preschoolers think biological properties are more likely to be passed from parent to child than are social or psychological properties. They argue that this implies that the children have a biology-like inheritance theory. The evidence concerning this issue is complex. On one hand, Solomon et al. (1996) claim that preschoolers do not have a biological concept of inheritance, because they do not have an adult's understanding of the biological causal mechanism involved. On the other hand, there is growing cross-cultural evidence that 4-5 year old children believe (like adults) that the identity of animals and plants follows that of their progenitors, regardless of the environment in which the progeny matures (e.g., progeny of cows raised with pigs, acorns planted with apple seeds) (Gelman & Wellman 1991; Atran et al., 2001; Sousa et al., 2002). Furthermore, it appears that Carey's (1985) results on psychology versus biology may only hold for urban children who have little intimate contact with nature (Atran, et al., 2001; Ross et al., 2003). Altogether, the evidence suggests that 4-5 year old children do have a distinct biology, though perhaps one without a detailed model of causal mechanisms (See Rozenbilt and Keil, 2002, for evidence that adults also only have a superficial understanding of mechanisms).

Domains and brain regions. Are these hypothesized domains associated with dedicated brain structure? There is intriguing evidence concerning category-specific deficits where patients may lose their ability to recognize and name category members in a particular domain of concepts. For example, Nelson (1946) reported a patient who was unable to recognize a telephone, a hat, or a car but could identify people and other living things (the opposite pattern is also observed and is more common). These deficits are consistent with the idea that anatomically and functionally distinct systems represent living versus nonliving things (Sartori & Job, 1988). An alternative claim (e.g., Warrington & Shallice, 1984) is that these patterns of deficits are due to the fact that different kinds of information aid in categorizing different kinds of objects. For example, perceptual information may be relatively more important for recognizing living kinds and functional information more important for recognizing artifacts (see Farah & McClelland, 1991; Devlin et al., 1998 for computational implementations of these ideas). Although, the weight of evidence appears to favor the kinds of information view (see Damasio et al., 1996; Forde, in press; Forde & Humphreys, in press; Simmons & Barsalou, in press), the issue continues to be debated (see Caramazza & Shelton, 1998, for a strong defense of the domain specificity view).

Domains and memory. The issue of domain specificity returns us to one of earlier themes: Does memory organization depend on the meaning? We've seen that early research on semantic memory was problematic in this respect, since many of the findings that investigators used to support meaning-based organization had alternative explanations. General-purpose decision processes could produce the same pattern of results, even if the information they operated on was haphazardly organized. Of course, in those olden days, semantic memory was supposed to be a hierarchically organized network like that in Figure 1; the network clustered concepts through shared superordinates and properties but was otherwise undifferentiated. Modularity and domain specificity offer a new take on semantic-based memory structure—a partition of memory space into distinct theoretical domains. Can large-scale theories like these support memory organization in a more adequate fashion than homogeneous networks?

One difficulty in merging domain specificity with memory structure is that domain theories don't taxonomize categories, they taxonomize assumptions. What differentiates domains is the set of assumptions or warrants they make available for thinking and reasoning (see Toulmin, 1958, for one such theory), and this means that a particular category of objects usually falls in more than one domain. To put it another way, domain-specific theories are "stances" (Dennett, 1971) or "construals" (Keil, 1995) that overlap in their instances. Take the case of people. The naive psychology domain treats people as having beliefs and goals that lend themselves to predictions about actions (e.g., Leslie, 1987; Wellman, 1990). The naive physics domain treats people as having properties like mass and velocity that warrant predictions about support and motion (e.g., McCloskey, 1983; Clement, 1983). The naive law-school domain treats people as having properties, such as social rights and responsibilities, that lead to predictions about obedience or deviance (e.g., Fiddick, Cosmides, & Tooby, 2000). The naive biology domain (at least in the Western adult version) treats people as having properties like growth and self-animation that lead to expectations about behavior and development. In short, each ordinary category may belong to many domains.

If domains organize memory, then long-term memory will have to store a concept in each of the domains to which it is related. Such an approach makes some of the difficulties of the old semantic memory more perplexing. Recall the issue of identifying the same concept across individuals, which we discussed earlier (see *Concepts as Positions in Memory Structure*). Memory modules have the same problem, but they add to it the dilemma of identifying concepts *within* individuals. How do you know that PEOPLE in your psychology module is the same concept as PEOPLE in your physics module and PEOPLE in your law-school module? Similarity is out (since the modules won't organize them in the same way), spelling is out (both concepts might be tied to the word "people" in an internal dictionary, but then fungi and metal forms are both tied to the word "mold"), and interconnections are out (since they would defeat the idea that memory is organized by domain). We can't treat the multiple PEOPLE

concepts as independent either, since it's important to get back and forth between them. For example, the rights-and-responsibilities information about people in your law-school module has to get together with the goals-and-desires information about people in your psychology module in case you have to decide, together with your fellow jury members, whether the killing was a hate crime or was committed with malice aforethought.

It is reasonable to think that background theories provide premises or grounds for inferences about different topics, and it is also reasonable to think that these theories have their “proprietary concepts.” But if we take domain-specific modules as the basis for memory structure—as a new semantic memory—we have to worry about nonproprietary concepts, too. We've argued that there must be such concepts, since we can reason about the same thing with different theories. Multiple storage is a possibility, if you're willing to forego memory economy and parsimony, and if you can solve the identifiability problem that we discussed in the previous paragraph. Otherwise, these domain-independent concepts have to inhabit a memory space of their own, and modules can't be the whole story.

Summary. We seem to be arriving at a skeptical position with respect to the question of whether memory is semantically organized, but we need to be clear about what is and what is not in doubt. What we doubt is that there is compelling evidence that long-term memory is structured in a way that mirrors lexical structure, as in the original semantic-memory models. We don't doubt that memory reflects meaningful relations among concepts, and it is extremely plausible that these relations depend to some extent on word meanings. For example, there may well be a relation in memory that links the concept TRUCKER with the concept BEER, and the existence of this link is probably due in part to the meaning of “trucker” and “beer.” What is not so clear is whether memory structure directly reflects the sort of relations that, in linguistic theory, organizes the meaning of words (where, e.g., “trucker” and “beer” are probably not closely connected). We note, too, that we have not touched (and we don't take sides on) two related issues, which are themselves subjects of controversy.

One of these residual issues is whether there is a split in memory between (a) general knowledge and (b) personally experienced information that is local to time and place. *Semantic memory* (Tulving, 1972) or *generic memory* (Hintzman, 1978) is sometimes used as a synonym for general knowledge in this sense, and it is possible that memory is partitioned along the lines of this semantic/episodic difference, even though the semantic side is not organized by lexical content. The controversy in this case is how such a dual organization can handle learning of “semantic” information from “episodic” encounters (see Tulving, 1984, and his critics in the same issue of *Behavioral and Brain Sciences*, for the ins and outs of this debate).

The second issue that we’re shirking is whether distributed brands of connectionist models can provide a basis for meaning-based memory. One reason for shirking is that distributed organization means that concepts like DAISY and CUP are *not* stored according to their lexical content. Instead, parts of the content of each concept are smeared across memory in overlapping fashion. It’s possible, however, that at a subconcept level—at the level of features or hidden units—memory has a semantic dimension, and we must leave this question open.

Conclusions and Future Directions

Part of our charge was to make some projections about the future of research on concepts. We don’t recommend a solemn attitude toward our predictions. But there are several trends that we have identified and, barring unforeseen circumstances (never a safe assumption), these trends should continue. One property our nominations share is that they uniformly broaden the scope of research on concepts. Here’s our shortlist.

Sensitivity to multiple functions (see also Solomon et al., 1999). The prototypical categorization experiment involves training undergraduates for about an hour and then giving transfer tests to assess what they have learned. This practice is becoming increasingly atypical, even among researchers studying

artificially constructed categories in the lab. Recently researchers have studied functions other than categorization, as well as interactions across functions.

Broader applications of empirical generalizations and computational models. As a wider range of conceptual functions come under scrutiny, new generalizations emerge and computational models face new challenges (e.g., Yamauchi, et al, 2002). Both developments set the stage for better bridging to other contexts and applications. This is perhaps most evident in the area of cognitive neuroscience where computational models have enriched studies of multiple categorization and memory systems (and vice versa). Norman, Brooks, Coblenz, and Babcock (1992) provide a nice example of extensions from laboratory studies to medical diagnosis in the domain of dermatology.

Greater interactions between work on concepts and psycholinguistic research. We've pressed the point that research on concepts has diverged from psycholinguistics because two different concepts of concepts seem to be in play in these fields. But it can't be true that the concepts we use in online sentence understanding are unrelated to the concepts we employ in reasoning and categorizing. There is an opportunity for theorists and experimenters here to provide an account of the interface between these functions. One possibility, for example, is to use sentence comprehension techniques to track the way that the lexical content of a word in speech or text is transformed in deeper processing (see Pinango, Zurif, & Jackendoff, 1999, for one effort in this direction). Another type of effort at integration is Wolff and Song's (2003) work on causal verbs and people's perception of cause, where he contrasts predictions derived from cognitive linguistics with those from cognitive psychology.

Greater diversity of participant populations. Although research with USA undergraduates at major universities will probably never go out of style (precedent and convenience are two powerful staying forces), we expect the recent increase to continue in the use of other populations. Work by Nisbett and his associates (e.g. Nisbett, Peng, Choi, & Norenzayan, 2001; Nisbett & Norenzayan, 2002) has called into question the idea that basic cognitive processes are universal, and categories and conceptual

functions are basic cognitive functions. In much of the work by Atran, Medin and their associates, undergraduates are the “odd group out” in the sense that their results deviate from those of other groups. In addition, cross-linguistic studies are often an effective research tool for addressing questions about the relationship between linguistic and conceptual development (e.g., Waxman, 1999).

More psychometaphysics. An early critique of the theory theory is that it suffered from vagueness and imprecision. As we’ve seen in this review, however, this framework has led to more specific claims (e.g. Ahn’s causal status hypothesis) and the positions are clear enough to generate theoretical controversies (e.g. contrast Smith, Jones, & Landau, 1996 with Gelman, 2000, and Booth & Waxman, 2002, in press, with Smith, Jones, Yoshida, & Colunga, 2003). It is safe to predict even greater future interest in these questions.

All of the above in combination. Concepts and categories are shared by all the cognitive sciences, so there’s very little room for researchers to stake out a single paradigm or subtopic and work in blissful isolation. Although the idea of a semantic memory uniting memory structure, lexical organization, and categorization may have been illusory, this doesn’t mean that progress is possible by ignoring the insights on concepts that these perspectives (and others) provide. We may see further fragmentation in the concepts of concepts, but it will still be necessary to explore the relations among them. Our only firm prediction is that the work we will find most exciting will be research that draws on multiple points of view.

References

- Ahn, W-K. (1998). Why are different features central for natural kinds and artifacts?: The role of causal status in determining feature centrality. Cognition, 69, 135-178.
- Ahn, W-K., Kalish, C., Gelman, S. A., Medin, D. L., Luhmann, C., Atran, S., et al (2001). Why essences are essential in the psychology of concepts. Cognition, 82, 59-69.
- Anderson, J. R. (1990). The adaptive character of thought. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1991). Is human cognition adaptive? Behavioral and Brain Sciences, 14, 471-517.
- Anderson, J. R., & Bower, G. H. (1973). Human associative memory. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. & Fincham, J. M. (1996). Categorization and sensitivity to correlation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 259-277.
- Ashby, F.G., Alfonso-Reese, L.A., Turken, A.U. & Waldron, E.M. (1998). A neuropsychological theory of multiple systems in category learning. Psychological Review, 105, 442-481.
- Ashby, F.G., & Maddox, W.T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. Journal of Mathematical Psychology, 37, 372-400.
- Atran, S. (1990). Cognitive foundations of natural history. Cambridge, UK: Cambridge University Press.
- Atran, S. (1999). Itzaj Maya folk-biological taxonomy. In D. Medin & S. Atran (eds.), Folk biology. Cambridge MA: MIT Press.
- Atran, S., Medin, D., Lynch, E., Vapnarsky, V., Ucan Ek', E. & Sousa, P. (2001). Folkbiology doesn't come from folkpsychology: Evidence from Yukatek Maya in cross-cultural perspective. Journal of Cognition and Culture 1:3-42.
- Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. Behavioral and Brain Sciences, 21, 547-609.
- Atran, S. & Sperber, D. (1991). Learning without teaching: Its place in culture. In L. Tolchinsky-Landsmann (ed.), Culture, schooling and psychological development. Norwood NJ: Ablex.
- Atran, S. (1985). The nature of folk-botanical life forms. American Anthropologist, 87, 298-315

Bailenson, J.N., Shum, M., Atran, S., Medin, D.L. & Coley, J.D.(2002). A Bird's eye View: Biological Categorization and Reasoning Within and Across Cultures. Cognition, 84, 1-53.

Balota, D. A. (1994). Visual word recognition: A journey from features to meaning. In M. A. Gernsbacher (Ed.), Handbook of psycholinguistics (pp. 303-358). San Diego, CA: Academic Press.

Barsalou, L. W. (1983). Ad-hoc categories. Memory and Cognition, 11, 211-227.

Blok, S., Newman, G., & Rips, L. J. (in press). Individuals and their concepts. In W-k. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the lab*. Washington, D.C.: American Psychological Association.

Bourne, L.E. Jr. (1970). Knowing and using concepts. Psychological Review, 77,546-556

Brooks, L.R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B.B. Lloyd (Eds.), Cognition and Categorization(pp.169-211) New York:Wiley

Bruner, J.S., Goodnow, J.J. & Austin, G.A. (1996). A Study of Thinking. New York: Wiley

Burge, T. (1999). Comprehension and interpretation. In L. E. Hahn (Ed.), The Philosophy of Donald Davidson (pp. 229-250). Chicago: Open Court.

Busemeyer, J. R., Dewey, G. I., & Medin, D. L. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 638-648.

Caramazza, A., & Grober, E. (1976). Polysemy and the structure of the subjective lexicon. Georgetown University Round Table on Language and Linguistics, 181-206.

Carey, S. (1985). Conceptual change in childhood. Cambridge, MA: MIT Press.

Carey, S. (1995). On the origin of causal understanding. In D. Sperber, D. Premack & A. J. Premack (Eds.), Causal cognition: A multidisciplinary debate (268-308). New York: Oxford University Press.

Chierchia, G., & McConnell-Ginet, S. (1990). Meaning and grammar: An introduction to semantics. Cambridge, MA: MIT Press.

Clapper, J. & Bower, G. (2002). Adaptive categorization in unsupervised learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 28(5)

Clement, J. (1983). A conceptual model discussed by Galileo and used intuitively by physics students. In D. Gentner & A. L. Stevens (Eds.), Mental models (pp. 325-340). Hillsdale, NJ: Erlbaum.

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

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8, 240-247.

Conrad, F. G., & Rips, L. J. (1986). Conceptual combination and the given/new distinction. Journal of Memory and Language, 25, 255-278.

Dennett, D. C. (1971). Intensional systems. Journal of Philosophy, 68, 87-106.

Erickson, M.A. & Kruschke, J.K. (1998). Rules and exemplars in category learning. Journal of Experimental Psychology: General, 127, 107-140

Fiddick, L., Cosmides, L., & Tooby, J. (2000). No interpretation without representation: The role of domain-specific representations and inferences in the Wason selection task. Cognition, 77, 1-79.

Fillmore, C. J., & Atkins, B. T. S. (2000). Describing polysemy: The case of 'crawl.' In Y. Ravin & C. Leacock (Eds.), Polysemy: Theoretical and computational approaches (pp. 91-110). Oxford, England: Oxford University Press.

Filoteo, J.V., Maddox, W.T., & Davis, J.D. (2001) A possible role of the striatum in linear and nonlinear categorization rule learning: Evidence from patients with Huntington's disease. Behavioral Neuroscience, 115, 786-798.

Fodor, J. (1994). Concepts: A potboiler. Cognition, 50, 95-113.

Fodor, J. (1998). Concepts: Where cognitive science went wrong. Oxford, England: Oxford University Press.

Franks, B. (1995). Sense generation: A "quasi-classical" approach to concepts and concept combination. Cognitive Science, 19, 441-505.

Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-head combinations. Journal of Experimental Psychology: Learning, Memory, & Cognition, 23, 71-87.

Gelman, S.A. (2003). The essential child: Origins of essentialism in everyday thought. Oxford, UK: Oxford University Press.

Gelman, S.A. & Coley, J.D. (1990). The importance of knowing a dodo is a bird: Categories and inferences in 2-year-old children. Developmental Psychology, 26(5), 796-804.

Gelman, S.A. & Koenig, M.A. (2001). The role of animacy in children's understanding of "move." Journal of Child Language, 28(3), 683-701.

Gelman, S. A. (2000). The role of essentialism in children's concepts. In H. W. Reese (Ed.), Advances in child development and behavior, Vol. 27 (pp. 55-98). San Diego: Academic Press.

Gelman, S. A., & Hirschfeld, L. A. (1999). How biological is essentialism? In D. L. Medin & S. Atran (Eds.), Folkbiology (pp. 403-446). Cambridge, MA: MIT Press.

Gelman, S.A. & Wellman, H.M. (1991). Insides and essence: Early understandings of the non-obvious. Cognition, 38(3), 213-244.

Gelman, S. A., Star, J. R., & Flukes, J. E. (2002). Children's use of generics in inductive inference. Journal of Cognition and Development, 3, 179-199.

Gleitman, L. R., & Gleitman, H. (1970). Phrase and paraphrase. New York: W.W. Norton.

Goldstone, R.L. (1998). Perceptual learning. Annual Review of Psychology, 49, 585-612.

Goldstone, R. L. (2003). Learning to perceive while perceiving to learn. In R. Kimchi, M. Behrmann, & C. Olson (Eds.), Perceptual organization in vision: Behavioral and neural perspectives (pp. 233-278). Mahwah, NJ: Erlbaum.

Goldstone, R.L., Lippa, Y., & Shiffrin, R.M. (2001). Altering object representations through category learning. Cognition, 78, 27-43.

Goldstone, R.L., & Rogosky, B.J. (2002). Using relations within conceptual systems to translate across conceptual systems. Cognition, 84, 295-320.

Goldstone, R.L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. Journal of Experimental Psychology: General, 130, 116-139.

Gutheil, G., & Rosengren, K. S. (1996). A rose by any other name: preschoolers understanding of individual identity across name and appearance changes. British Journal of Developmental Psychology, 14, 477-498.

- Hampton, J. (1979). Polymorphous concepts in semantic memory. Journal of Verbal Learning and Verbal Behavior, 18, 441-461.
- Hampton, J. (1987). Inheritance of attributes in natural concept conjunctions. Memory & Cognition, 15, 55-71.
- Hampton, J. (1997). Conceptual combination. In K. Lamberts & D. Shanks (Eds.), Knowledge, concepts, and categories (pp. 133-159). Cambridge, MA: MIT Press.
- Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. Bulletin of the Psychonomic Society, 28, 242-247
- Hintzman, D. L. (1978). The psychology of learning and memory. San Francisco: Freeman.
- Hintzman, D. L. (1986). 'Schema abstraction' in a multiple-trace memory model. Psychological Review, 93, 411-428.
- Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. Developmental Psychology, 38, 883-894.
- Homa, D., Sterling, S. & Trepel, L. (1981). Limitations of exemplar based generalization and the abstraction of categorical information. Journal of Experimental Psychology: Human Learning and Memory, 7, 418-439.
- Johansen, M.J. & Palmeri, T.J. (in press). Are There Representational Shifts in Category Learning? Cognitive Psychology.
- Johnson, C., & Keil, F. (2000). Explanatory understanding and conceptual combination. In F. C. Keil & R. A. Wilson (Eds.), Explanation and cognition (pp. 327-359). Cambridge, MA: MIT Press.
- Keil, F.C. (1989). Concepts, Kinds, and Cognitive Development. Cambridge, MA: MIT Press.
- Keil, F. C. (1981). Constraints on knowledge and cognitive development. Psychological Review, 88(3), 197-227.
- Keil, F. C. The birth and nurturance of concepts by domains: The origins of concepts of living things. In Hirschfeld, L.A. & Gelman, S.A. (Ed). Mapping the mind: Domain specificity in cognition and culture. (pp. 234-254).
- Keil, F.C., Levin, D.T., Richman, B.A., & Gutheil, G. (1999). Mechanism and explanation in the development of biological thought: The case of disease. In D. Medin & S. Atran (eds.), Folkbiology (pp. 285-319). Cambridge, MA: MIT Press.

Keil, F. C. (1995). The growth of causal understanding of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), Causal cognition (pp. 234-262). Oxford, England: Oxford University Press.

Kellman, P.J. & Spelke, E.S. (1983). Perception of partly occluded objects in infancy. Cognitive Psychology, 15, 483-524.

Klein, D. E., & Murphy, G. L. (2002). Paper has been my ruin: conceptual relations of polysemous senses. Journal of Memory and Language, 47, 548-570.

Knapp, A.G. & Anderson, J.A. (1984). Theory of categorization based on distributed memory storage. Journal of Experimental Psychology: Learning, Memory and Cognition, 10, 616-637.

Knowlton, B.J. & Squire, L.R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. Science, 262, 1747-1749.

Knowlton, B.J., Mangels, J.A. & Squire, L.R. (1996). A neostriatal habit learning system in humans. Science, 273, 1399-1402.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Link, G., & Chierchia, G. (1995). Genericity: An Introduction. In G. N. Carlson & F. J. Pelletier (Eds.), The generic book (pp. 1-124). Chicago: University of Chicago Press.

Kruschke, J.K. (1992).ALCOVE: An exemplar based connectionist model of category learning. Psychological Review, 99, 22-44.

Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. Cognitive Science, 14, 551-577.

Lamberts, K. (1995). Categorization under time pressure. Journal of Experimental Psychology: General 124, 161-180.

Landauer, T. K, & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychological Review, 104, 211-240.

Leslie, A. M. (1987). Pretense and representation: The origins of “theory of mind.” Psychological Review, 94, 412-426.

Liittschwager, J. C. (1995). Children’s reasoning about identity across transformations. Dissertation Abstracts International, 55 (10), 4623B. (UMI No. 9508399).

Love, B. C., Markman, A. B., & Yamauchi, T. (2000). Modeling inference and classification learning. The National Conference on Artificial Intelligence (AAAI-2000), 136-141.

Love, B.C., Medin, D.L, & Gureckis, T.M (in press). SUSTAIN: A network model of category learning. Psychological Review.

Lucas, M. (2000). Semantic priming without association. Psychonomic Bulletin & Review, 7, 618-630.

Lyons, J. (1977). Semantics (vol. 2). Cambridge, England: Cambridge University Press.

Maddox, W.T. (1999). On the dangers of averaging across observers when comparing decision bound models and generalized context models of categorization. Perception and Psychophysics, 61, 354-375.

Maddox, W.T. (2002) Learning and attention in multidimensional identification, and categorization: separating low-level perceptual processes and high level decisional processes. Journal of Experimental Psychology: Learning, Memory, & Cognition, 28, 99-115.

Maddox, W.T., & Ashby, F.G. (1993). Comparing decision bound and exemplar models of categorization. Perception & Psychophysics, 53, 49-70

Maddox, W.T. & Ashby, F.G. (1998). Selective attention and the formation of linear decision boundaries: Comment on McKinley and Nosofsky (1996). Journal of Experimental Psychology: Human Perception and Performance, 24, 301-321.

Malt, B.C. (1994). Water is not H-sub-2O. Cognitive Psychology, 27(1), 41-70.

Malt, B. C., Ross, B.H. & Murphy, G.L. (1995). Predicting features for members of natural categories when categorization is uncertain. Journal of Experimental Psychology: Learning, Memory and Cognition, 21,646-661.

Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing vs. naming: Similarity and the linguistic categorization of artifacts. Journal of Memory and Language, 40, 230-262

Malt, B. C. & Smith, E. E. (1984). Correlated properties in natural categories. Journal of Verbal Learning & Verbal Behavior, 2: 250-269.

Markman, A.B. & Makin, V.S.(1998).Referential communication and category acquisition. Journal of Experimental Psychology:General,127,331-354.

McCloskey, M., & Glucksberg, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. Cognitive Psychology, *11*, 1-37.

McKinley, S.C. & Nosofsky, R.M. (1995). Investigations of exemplar and decision bound models in large, ill defined category structures. Journal of Experimental Psychology: Human Perception and Performance, *21*, 128-148.

Medin, D. (1989) Concepts and conceptual structures. American Psychologist *45*, 1469-1481.

Medin, D.L., Ross, N., Atran, S., Burnett, R.C., & Blok, S. V. (2002) Categorization and reasoning in relation to culture and expertise. In Ross, B. H. (Ed). The psychology of learning and motivation: Advances in research and theory, *41*, 1-41.

Medin, D. L. (1986). Commentary on "Memory Storage and Retrieval Processes in Category Learning." Journal of Experimental Psychology: General, *115*, 373-381.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, *8*, 37-50.

Medin, D. L. & Coley, J. D. (1998). Concepts and categorization. In J. Hochberg (Ed.), Handbook of perception and cognition. Perception and cognition at century's end: History, philosophy, theory (pp. 403-439). San Diego: Academic Press.

Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? Annual Review of Psychology, *51*, 121-147.

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

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou and A. Ortony (Eds.), Similarity and analogical reasoning (pp. 179-195). New York: Cambridge University Press.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. Cognitive Psychology, *20*, 158-190.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. Journal of Experimental Psychology, *90*, 227-234.

Murphy, G. L. (1988). Comprehending complex concepts. Cognitive Science, *12*, 529-562.

- Murphy, G. L. (2002). The big book of concepts. Cambridge, MA: MIT Press.
- Murphy, G. L., & Ross, B.H. (1994). Predictions from uncertain categorizations. Cognitive Psychology, 27, 148-193
- Nisbett, R., Peng, K., Choi, I. & Norenzayan, A. (2001) Culture and systems of thought: Holistic vs. analytic cognition. Psychological Review, 108, 291-310.
- Nisbett, R.E. & Norenzayan, A. (2002). Culture and cognition. In Pashler, H. & Medin, D. (Eds.), Stevens' Handbook of Experimental Psychology (3rd ed.), Vol. 2: Memory and cognitive processes, 561-597.
- Norman, D. A., & Rumelhart, D. E. (1975). Explorations in cognition. San Francisco: W. H. Freeman.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, 115, 39-57.
- Nosofsky, R. M. (1991). Tests of an exemplar model for relation perceptual classification and recognition in memory. Journal of Experimental Psychology: Human Perception and Performance, 17, 3-27
- Nosofsky, R.M. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. Psychological Science, 9, 247-255.
- Nosofsky, R. M. & Johansen, M.K. (2000). Exemplar based accounts of "multiple system" phenomena in perceptual categorization. Psychonomic Bulletin and Review, 7, 375-402.
- Nosofsky, R. M. & Palmeri, T.J. (1997a). An exemplar based random walk model of speeded classification. Psychological Review, 104, 266-300.
- Nosofsky, R. M. & Palmeri, T.J. (1997b). Comparing exemplar retrieval and decision-bound models of speeded perceptual classification. Perception and Psychophysics, 59, 1027-1048.
- Nosofsky, R. M., Palmeri, T.J. & McKinley, S.C. (1994). Rule-plus-exception model of classification learning. Psychological Review, 101, 53-79.
- Nosofsky, R. M. & Zaki, S.R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar based interpretation. Psychological Science, 9, 247-255.

Nosofsky, R.M. & Zaki, S.R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. Journal of Experimental Psychology: Learning, Memory, & Cognition, 28, 924-940.

Nosofsky, R.M., Clark, S.E., & Shin, H.J. (1989). Rules and exemplars in categorization, identification, and recognition. Journal of Experimental Psychology: Learning, Memory, & Cognition, 15, 282-304.

Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. Cognition, 11, 35-58.

Palmeri, T.J. (1997). Exemplar similarity and the development of automaticity. Journal of Experimental Psychology: Learning, Memory and Cognition, 23, 324-354.

Palmeri, T.J. (1999). Learning hierarchically structured categories: A comparison of category learning models. Psychonomic Bulletin and Review, 6, 495-503.

Palmeri, T.J., & Flanery, M.A. (1999). Learning about categories in the absence of training: Profound amnesia and the relationship between perceptual categorization and recognition memory. Psychological Science, 10, 526-530.

Palmeri, T.J., & Flanery, M.A. (2002). Memory systems and perceptual categorization. In B.H. Ross (Ed.), The Psychology of Learning and Motivation (Vol. 41), Academic Press.

Partee, B. H. (1995). Lexical semantics and compositionality. In L. R. Gleitman & M. Liberman (vol. eds.) & D. N. Osherson (series ed.), Invitation to cognitive science (vol. 1: Language). Cambridge, MA: MIT Press.

Pinango, M. M., Zurif, E., & Jackendoff, R. (1999). Real-time processing implications of enriched composition at the syntax-semantics interface. Journal of Psycholinguistics Research, 28, 395-414

Posner, M.I. & Keele, S.W.(1968). On the genesis of abstract ideas. Journal of Experimental Psychology,77,353-363.

Posner, M.I. & Keele, S.W.(1970).Retention of abstract ideas. Journal of Experimental Psychology,83,304-308.

Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. Behavioral Sciences, 12, 410-430.

Quillian, M. R. (1969). The teachable language comprehender: A simulation program and theory of language. Communications of the ACM, 12, 459-476.

- Reagher, G. & Brooks, L.R. (1993). Perceptual manifestations of an analytic structure: the priority of holistic individuation. Journal of Experimental Psychology: General, 122, 92-114
- Reber, P.J., Stark, C.E.L, & Squire, L.R. (1998a). Cortical areas supporting category learning identified using functional MRI. Proceedings of the National Academy of Sciences of the USA, 95, 747-750
- Reber, P.J., Stark, C.E.L, & Squire, L.R. (1998b). Contrasting cortical activity associated with category memory and recognition memory. Learning and Memory, 5, 420-428.
- Reed, S.K. (1972). Pattern recognition and categorization. Cognitive Psychology, 3, 382-407.
- Rehder, B., & Hastie, R. (2001). The essence of categories: The effects of underlying causal mechanisms on induction, categorization, and similarity. Journal of Experimental Psychology: General, 130, 323-360.
- Restle, F. (1962). The selection of strategies in cue learning. Psychological Review, 69, 329-343.
- Rips, L. J. (1995). The current status of research on concept combination. Mind & Language, 10, 72-104.
- Rips, L. J. (2001). Necessity and natural categories. Psychological Bulletin, 127, 827-852.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. Journal of Verbal Learning and Verbal Behavior, 12, 1-20.
- Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), Cognitive development and the acquisition of language (pp. 111-144). New York: Academic Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), Cognition and categorization (pp. 27-48). Hillsdale, NJ: Erlbaum.
- Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. Cognitive Psychology, 7, 573-605.
- Rosch E, Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976). Basic objects in natural categories. Cognitive Psychology, 8, 382-439.
- Rozenblit, L., and Keil, F. (2002). The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth. Cognitive Science, 26, 521-562.

- Ross, B.H. (1997). The use of categories affects classification. Journal of Memory and Language, 37, 240-267.
- Ross, B.H. (1999). Postclassification category use: The effects of learning to use categories after learning to classify. Journal of Experimental Psychology: Learning, Memory and Cognition, 25, 743-757.
- Ross, B.H. (2000). The effects of category use on learned categories. Memory & Cognition, 28, 51-63.
- Ross, B.H. & Murphy, G.L. (1996). Category based predictions: Influence of uncertainty and feature associations. Journal of Experimental Psychology: Learning, Memory and Cognition, 22, 736-753.
- Schyns, P., Goldstone, R. & Thibaut, J. (1998). Development of features in object concepts. Behavioral and Brain Sciences 21, 1-54.
- Schyns, P. & Rodet, L. (1997). Categorization creates functional features. Journal of Experimental Psychology: Learning, Memory and Cognition.23, 681-696
- Slooman, S.A. (1993). Feature-based induction. Cognitive Psychology 25, 231-280.
- Slooman, S. A., & Malt, B. (in press). Artifacts are not ascribed essences, nor are they treated as belonging to kinds. Language and Cognitive Processes.
- Slooman, S.A., Love, B.C., & Ahn, W.-K. (1998). Feature centrality and conceptual coherence. Cognitive Science, 22(2), 189-228.
- Smith, L.B., Jones, S. S., & Landau, B. (1996). Naming in young children: A dumb attentional mechanism? Cognition, 60(2), 143-171.
- Smith, L.B., Jones, S.S. Yoshida, H., & Colunga, E. (2003). Whose DAM account? Attentional learning explains Booth and Waxman. Cognition, 87(3), 209-213.
- 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, 485-527.
- Smith, E. E., Shoben, E.J. & Rips, L.J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. Psychological Review.81,214-241.

Smith, J.D. & Minda, J.P. (1998). Prototypes in the mist: The early epochs of category learning. Journal of Experimental Psychology: Learning, Memory and Cognition, 24, 1411-1436

Smith, J.D. & Minda, J.P. (2000). Thirty categorization results in search of a model. Journal of Experimental Psychology: Learning, Memory and Cognition, 26, 3-27

Smith, J.D., Murray, M.J., Jr. & Minda, J.P. (1997). Straight talk about linear separability. Journal of Experimental Psychology: Learning, Memory and Cognition, 23, 659-680

Sober, E. (1980). Evolution, population thinking, and essentialism. Philosophy of Science, 47, 350-383.

Solomon, K. O., Medin, D. L., & Lynch, E. B. (1999). Concepts do more than categorize. Trends in Cognitive Science, 3, 99-105.

Stanton, R., Nosofsky, R.M. & Zaki, S. (2002). Comparisons between exemplar similarity and mixed prototype models using a linearly separable category structure. Memory & Cognition, 30,

Storms, G., de Boeck, P., van Mechelen, I., & Ruts, W. (1998). No guppies, nor goldfish, but tumble dryers, Noriega, Jesse Jackson, panties, car crashes, bird books, and Stevie Wonder. Memory & Cognition, 26, 143-145.

Stevens, M. (2000). The essentialist aspect of naïve theories. Cognition, 74, 149-175.

Stevens, M. (2001). Only causation matters. Cognition, 82, 71-76.

Toulmin, S. (1958). The uses of argument. Cambridge, England: Cambridge University Press.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving and W. Donaldson (Eds.), Organization of memory (pp. 381-403). New York: Academic Press.

Tulving, E. (1984). Precis of Elements of episodic memory. Behavioral & Brain Sciences, 7, 223-268.

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

Verguts, T., Storms, G., & Tuerlinckx, F. (2001). Decision bound theory and the influence of familiarity. Psychonomic Bulletin and Review

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

Wellman, H.M., & Gelman, S.A. (1992) Cognitive development: Foundational theories of core domains. Annual Review of Psychology, 43, 337-375.

- Wiggins, D. (1980). Sameness and substance. Cambridge, MA: Harvard University Press.
- Wilcox, T., & Baillargeon, R. (1998). Object individuation in infancy: The use of featural information in reasoning about occlusion events. Cognitive Psychology, 37, 97-155.
- Wisniewski, E. J. (1997). When concepts combine. Psychonomic Bulletin & Review, 4, 167-183.
- Wisniewski, E. J. (2002). Concepts and categorization. In D.L. Medin (Ed.) Steven's Handbook of Experimental Psychology (3rd ed., pp. 467-532). Wiley: NewYork.
- Wisniewski, E. J. & Medin, D. L. (1994). On the interaction of theory and data in concept learning. Cognitive Science, 18, 221-281.
- Wolff, P., & Song, G. (2003). Models of causation and the semantics of causal verbs. Cognitive Psychology, 47, 241-275.
- Wolff, P., Medin, D., & Pankratz, C. (1999). Evolution and devolution of folkbiological knowledge. Cognition, 73, 177-204.
- Xu, F. (In press). The development of object individuation in infancy. In J. Fagen & H. Hayne (Eds.), Progress in infancy research (Vol. 3). Mahwah, NJ: Erlbaum.
- Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. Cognitive Psychology, 30, 111-153.
- Yamauchi, T., Love, B. C., & Markman, A. B. (2002). Learning non-linearly separable categories by inference and classification. Journal of Experimental Psychology: Learning, Memory & Cognition, 28 (3), 585-593.
- Yamauchi, T. & Markman, A.B. (1998). Category learning by inference and classification. Journal of Memory and Language, 39, 124-149.
- Yamauchi, T. & Markman, A.B. (2000a). Inference using categories. Journal of Experimental Psychology: Learning, Memory and Cognition, 26, 776-795.
- Yamauchi, T. & Markman, A.B. (2000b). Learning categories composed of varying instances: The effect of classification, inference and structural alignment. Memory & Cognition, 28, 64-78.
- Zadeh, L. (1965). Fuzzy sets. Information and Control, 8, 338-353.

Table 1. Some Theories of Concept Combination

Model	Domain	Representation of Head Noun	Modification Process
Hampton (1987)	Noun-Noun and Noun-Relative-Clause NPs (conjunctive NPs, e.g., <i>sports that are also games</i>)	Schemas (attribute-value lists with attributes varying in importance)	Modifier and head contribute values to combination on the basis of importance and centrality
Smith, Osherson, Rips, & Keane (1988)	Simple Adjective-Noun NPs (e.g., <i>red apple</i>)	Schemas (attribute-value lists with distributions of values and weighted attributes)	Adjective shifts value on relevant attribute in head and increases weight on relevant dimension.
Murphy (1988)	Adj-Noun and Noun-Noun NPs (esp. non-predicating NPs, e.g., <i>corporate lawyer</i>)	Schemas (lists of slots and fillers)	Modifier fills relevant slot; then representation is “cleaned up” on the basis of world knowledge.
Franks (1995)	Adj-Noun and Noun-Noun NPs (esp. privatives, e.g., <i>fake gun</i>)	Schemas (attribute-value structures with default values for some attributes)	Attribute-values of modifier and head are summed, with modifier potentially overriding or negating head values.
Gagné & Shoben (1997)	Noun-Noun NPs	Lexical representations containing distributions of relations in which nouns figure	Nouns are bound as arguments to relations (e.g., <i>flu virus</i> = virus causing flu).
Wisniewski (1997)	Noun-Noun NPs	Schemas (lists of slots and fillers, including roles in relevant events)	<ol style="list-style-type: none"> 1. Modifier noun is bound to role in head noun (e.g., <i>truck soap</i> = soap for cleaning trucks). 2. Modifier value is reconstructed in head noun (e.g., <i>zebra clam</i> = clam with stripes). 3. Hybridization (e.g., <i>robin canary</i> = cross between robin and canary)

Concepts and Categorization / 61

