

Emerging Ideas About Categories

Linda B. Smith
Indiana University

INTRODUCTION

The study of human categorization is contentious because it matters. How we understand the processes that give rise to categories is at the very core of how we understand human cognition. Over the past several decades, there has been growing concern that the assumptions that define and distinguish competing theories may not be quite right. These assumptions are based on the traditional metaphor that views categories as discrete, bounded things that are stable over time and context. In this view, categories are enduringly real, object-like, truly out there in the world and also in our heads. Thus, theorists in this tradition write about categories being *acquired*, *discovered*, and *possessed*. The boundedness and stability expected of categories is well exemplified in the following quote from Keil (1994):

Shared mental structures are assumed to be constant across repeated categorizations of the same set of instances and different from other categorizations. When I think about the category of dogs, a specific mental representation is assumed to be responsible for that category and roughly the same representation for a later categorization of dogs by myself or by another. (p. 169)

The problem with these traditional ideas is the fluidity of human categories that appear to be exquisitely malleable, adapting to fit the idiosyncrasies of the moment (e.g., Barsalou, 1993a, 1993b; Bransford & Johnson,

1972; Malt, 1994). We do not, it turns out, think *exactly* the same thing each time we think about a dog. Moreover, if you and I have different histories, we may well—even in the same context—think differently about the very same dog (e.g., Yoshida & Smith, 2003a, 2003b).

The chapters in this volume have been written in the midst of shifting assumptions about the fundamental nature of categories. There is not yet a well-accepted new framework, never mind a new theory, but research on these issues is clearly in flux. The goal of this commentary is to provide perspective on emerging ideas as represented in this volume.

A Brief History of Categories

The traditional treatment of categories derives from logic. This conceptualization, the logic of classes, distinguishes between the *extension* of a class and the *intension* of a class. The extension is all the possible members of a class. Thus the extension of the class “triangle” is all possible triangles. The intension is the rule that picks out all and only members of the class, for example, the intensional definition of a triangle might be a “closed figure having three sides.” Traditional psychological theories of categorization reflect these ideas of fixed extensions and intensions. In psychology, the extension is the “repeated categorizations” that people make, that is, the data to be explained. The intension is the hypothesized concept that determines the extension, the mental structure that proposed to cause people to categorize the way they do. Although this view has dominated theorizing and research on categories for nearly half a century, its theoretical problems have been evident for some time.

In the 1960s, theories of categorization attempted to explain the assumed fixed category extensions by internally represented intensive definitions that were lists of necessary and sufficient features. This approach came to be rejected on both theoretical and empirical grounds. First, there was no psychological basis for determining the features that form the primitives for concepts (Murphy & Medin, 1985). Second, no successful version of the theory was ever formulated—no one could find the defining properties of such everyday categories as dog or cow or game (see Rosch & Mervis, 1975; Smith & Medin, 1981). Third, there were data that directly contradicted the idea of necessary and sufficient features. Specifically, if a category is defined by necessary and sufficient features, all members should be equally good members. But the data say otherwise; people reliably judge some members of a category to be better than others. For example, a robin is a better example of a bird than is a penguin (e.g., Rosch, 1973).

In the 1970s, the field turned to probabilistic theories (Smith & Medin, 1981). These theories sought to explain human category judgments by ty-

ing them to general cognitive processes of memory, attention, association, and generalization by similarity. There are two versions, alike in the processes they assume, but different in their internal representations. And indeed, it has been argued that these two versions may not be formally distinguishable (Estes, 1986). By one account, known as prototype theory, concepts became lists of characteristic rather than defining features—a move that readily elevated robins over penguins as examples of birds (see Smith & Medin, 1981). By another, more radical version known as exemplar theory, concepts do not really exist in the sense of intensional definitions that determine category membership. Instead, people remember instances and associated properties (including associated language), and then general processes of memory retrieval, association, and generalization by similarity give rise to the in-task category judgments (e.g., Nosofsky, 1984; see also Smith & Medin, 1981). These exemplar accounts readily explain typicality effects and other effects suggesting probabilistic category membership. These theories also explain a wide array of experimental results on category learning, recognition, recall, and generalization (e.g., Nosofsky, 1984; Zaki & Nosofsky, 2001). Probabilistic feature-based theories also have had considerable recent success in modeling adults' judgments of common categories and their organization into domains (McRae, Cree, Westmacott, & de Sa, 1999; Vigliocco, Vinson, Lewis, & Garrett, 2004). In these efforts, the statistical regularities in large normative studies of everyday object and action categories have been shown to closely predict adult category and semantic judgments in a variety of tasks. Normative studies of the statistical structure of the first 300 nouns learned by children have also been shown to be predictive of their category learning (Samuelson & Smith, 2000; Smith, Colunga, & Yoshida, 2003).

However, probabilistic theories—both prototype and exemplar—have difficulty accounting for how people reason about categories. Specifically, people sometimes make category judgments that are decidedly more in accord with a defining feature view than a probabilistic view. For example, people will maintain that an organism that has no properties anything at all like a bird other than bird DNA and bird parents is, nonetheless, a bird (e.g., Rips, 1989; Keil, 1994).

In light of these last results, a number of researchers in the 1980s and 1990s turned to theory-like accounts of categories (e.g., Murphy & Medin, 1985). The idea here is that the mental structures that determine categories are naïve theories about the causal relatedness of different kinds of properties, both observable and nonobservable. For example, such a theory might include the following: Birds have wings and are lightweight *because* they fly and these behavioral and physical properties arise *because* of the genetic structure of birds. Accordingly, many researchers began studying people's beliefs about “really makes something what it is” and their reasoning

about the causal relatedness of properties relevant to category membership. The results of these studies suggest a distinction between core characteristics of things (and often not directly observable properties such as DNA) and the surface characteristics of things (for example, being bird shaped). That is, within intuitive theories some features are more important, and have more causal force, than others.

One version of the intuitive-theory account posits that people's theories about kinds are "essentialist" (see Gelman, 2003). The idea here is that people believe there is "an essence" that determines whether or not an instance is a member of a category. By this account, the reason that an organism that looks and acts nothing like a bird might still be judged to "really be a bird" is because the subject believes the organism possesses the essential but nonobvious properties that are true of all and only birds. These essentialist ideas thus resurrect the criterial-property concepts of the 1960s and the idea that a believed intension (a belief in an essential property) determines the extension (the belief in what really is a bird). However, by the modern-day essentialist perspective, it is not that instances actually share these properties or that these essential properties are even useful in recognizing instances in the world, but rather beliefs in the existence of these essential properties govern how people *reason* about category members. Moreover, these beliefs are organized by theory-like representations that causally relate instances and their properties.

Much contemporary research is devoted to the study of intuitive theories and their development. The intuitive theory view of concepts has opened new fields of study about categories—including induction, conceptual combination, and causal reasoning (E. Smith, 1989; Medin, 1989; Keil, 2003). Research on intuitive theories has also led to interesting insights about how reasoning differs in different domains (e.g., for biological versus nonbiological kinds, see Gelman, 2003).

Still, the naïve-theory view has its own problems. First, and as Ahn and Luhmann discuss in their chapter (chap. 11, this volume), there is no consensus as to what a naïve theory is, the formal nature of the representations, or the kinds of knowledge included. In general, naïve-theory theories are not as well defined or formalized as the probabilistic-feature sort, making rigorous testing of predictions difficult. Second, naïve theories clearly do not explain the full range of data traditionally viewed as the province of a theory of categorization. Instead, certain phenomena (induction, conceptual change, conceptual combination, and judgments of causal relatedness) are singled out as theoretically more important than phenomena concerning the recognition of instances. Thus, naïve theory accounts do not explain how one knows a bird when one sees (or hears) one, nor do they explain why robins are judged to be psychologically better birds than penguins. Moreover, the fact that people readily make these judgments is seen

as pretty much irrelevant to the intuitive-theory account of human categories (e.g., Armstrong, Gleitman, & Gleitman, 1983).

Third, naïve theories may not explain the very data they take to be their core phenomena. Keil (chap. 13, this volume; also Rozenblit & Keil, 2002) presents compelling data that even adults' naïve theories are often explanatorily inadequate and often quite incoherent. People believe they understand phenomena with greater precision and coherence than they really do. Keil (2003) suggested that people have at best coarse, not-quite right, and gap-filled understandings of the causal structure of even basic things. He provided evidence that their seeming causally based reasoning in laboratory experiments may derive from a combination of quite sketchy knowledge along with task-specific information. That is, causal reasoning appears to be made in the moment, ad hoc and on-the-fly, much like the ad hoc categories described earlier by Barsalou (1983). These in-the-moment temporary creations enable people to reason well, in ways adapted to the specific task—despite real gaps in their knowledge about causal relatedness. These results place naïve theories in the domain of situated cognition and real-time processes. In so doing, they forewarn that any complete theory of categories will require a specification of *both* knowledge (features correlations, theories) *and* the general cognitive processes on which they must depend.

General or Special Processes?

The dispute between probabilistic-feature accounts and naïve-theory accounts is also a debate about the fundamental nature of cognition. Is it based on general processes of perception, memory, attention, association, and generalization by similarity, or does it require other kinds of mechanisms and, in particular, propositional representations necessary to coherent and causal theories about how features are related? The differing positions are related to a disagreement about the most relevant data for a theory of categorization. The signature markers of domain general processes are frequency, instance, similarity, and typicality effects. These effects are ubiquitous throughout human (and animal) cognition. Yet some aspects of human cognition—category induction, conceptual combination, conceptual change—have been argued to be relatively immune from such effects (e.g., Barrett, Abdi, Murphy, & Gallagher, 1993; Murphy & Allopenna, 1994; Sloman, 1997; Gelman & Koenig, 2003). Proponents of the naïve theory view argue that phenomena that do not show these characteristic patterns are theoretically the most important, precisely because they cannot be easily explained by general processes and demand a special explanation in terms of propositional representations. Proponents of the naïve theory view admit that people do use general frequency and similar-

ity-based processes to make some category decisions. Keil, for example, suggested that people use only general process solutions in “desperation” when the particular category decision lacks the “nurturance and support of beliefs and principles” (1994, p. 239). Thus on one side are naïve-theory theorists who contend that certain aspects of human categorization cannot be explained by domain general cognitive processes, and who maintain that these phenomena are therefore most critical to a theory of categorization. As a consequence, these researchers study for the most part how people verbally reason about categories.

On the other side are those who seek explanations in terms of general processes, the side that includes probabilistic feature theories. These researchers concentrate on the relation between category decisions and what they see as the foundational processes: memory, attention, perceptual learning, and similarity (e.g., Barsalou, 2003; Chater & Vitanyi, 2003; Hampton, 2001; Goldstone & Steyvers, 2001; Nosofsky & Palmeri, 1998). Accordingly, these researchers pick problems to study that do not overlap much with those studied by naïve-theory theorists, problems such as the perceptual properties critical to the recognition of instances, attention, and perceptual learning.

There have been, however, some attempts by naïve-theory researchers to directly pit the two approaches against each other. The goal is to support the naïve theory side by showing that people’s category judgments do not depend strongly on similarity relations, the idea being that if similarity can be ruled out then propositional or theory-like beliefs are supported. The method used in these studies has not gone without criticism (see Jones & Smith, 1993). The experiments in this genre often present subjects with verbal descriptions of truly bizarre scenarios: animals that are magically transformed from one kind to another or the birth of babies with the DNA of one species and the appearance of another (e.g., Rips, 1989; Keil, 1991; Gelman & Coley, 1990). These fantastic scenarios are necessary because instances of the same real-world categories typically share deep conceptual similarities *and also* many perceptual and associative similarities, making it difficult with more ordinary scenarios to show that similarity does not matter (e.g., Jones & Smith, 1993).

Percepts or Concepts

In the developmental literature, this battle is fought over the issue of whether children’s categories are based on perceptual similarities such as shape and features such as wheels or whether their categories are based on conceptual features such as “can be eaten” or “used to carry water.” As Ahn points out in her chapter, this is a distortion of the larger theoretical issue that is not really about the kind of features but about the nature of the rep-

representations (feature correlations and counts vs. causal propositional structures). Ahn also notes that the perceptual–conceptual distinction is misleading in the developmental literature in that many of the “conceptual” features studied by developmental researchers are directly perceivable. Alternatively, these conceptual features could derive from direct associations among words (see, e.g., McRae et al., 1999; Vigliocco et al., 2004; Landauer & Dumais, 1997).

The developmental question has focused on the issue of percepts versus concepts primarily because of the developmental data themselves. A hundred years of research in developmental psychology suggests that preschool children often base their decisions on the static and currently perceptually available properties whereas older children base theirs on remembered properties or inferences. The fact of this general developmental trend is incontestable. What is contested is what it means. One possibility is that children have fundamentally different kinds of concepts than adults; what Keil (1994) suggested are “pseudoconcepts” whereas adults have theory-like or “true” concepts. The idea that children have different, more “illogical” concepts is one with a long history in developmental psychology (Bruner, 1986; Piaget, 1970; Wohwill, 1967).

For proponents of the naïve theory view, acceptance of this hypothesis would mean that developmentally immature categories based on perceptual similarity would have to change into developmentally mature categories based on intuitive theories. As Keil (1991) noted this would require “that coherent [intuitive] theories be able to develop out of something like networks of associations, that interconnected sets of explanatory beliefs can rise out of nothing more than probabilistic tabulations of features and relations” (p. 246). Keil went on to conclude, “This notion falters when one recognizes that there are not persuasive accounts in any domain showing how this might occur.” Thus the idea of a qualitative shift in the very nature of concepts is rejected out-of-hand by contemporary proponents of the naïve-theory account. It is not a position widely supported by the other side either. (For a third position supporting the idea of potentially qualitative shifts in representation see Fischer, Kenny, & Pipp, 1990; Andrews & Halford, 2002; Karmiloff-Smith, 1999.)

A second possibility, then, is that children’s concepts are like adult concepts but that both are based on general processes of association and generalization by similarity. This is the possibility that adherents of the naïve theory view must defend against and it has a number of supporters. These challenges to the naïve theory view attempt a direct assault by showing that the core phenomena assumed to be “conceptual” and *unexplainable* by general psychological processes can, in fact, be explained by ordinary processes of perception, attention, and memory. Attacks of this kind are now rising on many fronts, for example, in powerful similarity-based models (e.g.,

Hampton, 2001), in statistical learners fed only text as input (Landauer & Dumais, 1997), in successful feature-based models of common categories (McRae et al., 1999; Vigliocco et al., 2004), in Bayesian models that explain conceptual coherence essentially through feature tabulation and statistical inference (Tenenbaum & Griffiths, 2001), in connectionist models of category development (Smith et al., 2003), in explanations of cross-linguistic differences in categories (Yoshida & Smith, 2003a, 2003b; Sera et al., 2002), and in studies of perceptual learning and perceptual symbol systems (Goldstone & Barsalou, 1998). Rogers and McClelland (chap. 14, this volume) take a particularly comprehensive approach. They systematically demonstrate that connectionist networks can mimic many of the phenomena believed to be diagnostic of naïve theories and thus *not* based on general processes. Connectionist networks are associative learners that generalize by similarity and produce graded, context dependent decisions. This challenge thus goes to the core premise behind the segregation of human categorization data into two separate kinds: those explainable by general processes and those said not to be so explainable.

The third possibility is that children's concepts are like adult concepts, and are fundamentally at their core propositional and theory-like. By this hypothesis, children may often make judgments using the surface similarities of things, but there are aspects of their reasoning about categories that cannot be explained by associative learning, the tabulation of instances, and perceptual similarity. This is the specific possibility adherents to the naïve-theory view must support. The evidence for this position is, again, category judgments that do not seem to depend on the immediate perceptual input and that seem not to be influenced by similarity effects or that involve causal relations assumed to be unexplainable by general cognitive processes (e.g., Sloman, 1997; Gelman, 2003).

Is It Resolvable?

How can one decide between these two opposing views? Progress on this question appears stalled. This may be because the logical structure of the argument is itself flawed. Consider the structure of the two opposing claims.

General Process. The claim here is that one set of processes (which we can call G for general) explains human categories.

Naïve Theories. The claim here is only that there are *some* phenomena not explainable by G but that require other special mechanisms and representations, or S.

Three discouraging consequences arise in this structure.

1. It is a necessary truth that naïve-theory theorists can explain everything that general process accounts can. Anything explainable by G is necessarily explainable by G + S. This means naïve theory accounts are more powerful whereas general process accounts are more constrained and more parsimonious. It's a matter of "I can do more than you" versus "I can do (almost) everything you do, more elegantly."
2. The critical evidence for distinguishing these two approaches is a negative, showing that there are at least some aspects of human categorization that cannot be explained by general processes. This claim can never be disproven.
3. There is no way out of the debate. Each phenomenon that a naïve-theory theorist points to as special has some potential of being explained by a more powerful model of general processes (as in Rogers & McClelland, chap. 14, this volume). But for every phenomenon so explained there are other, perhaps even more special phenomena to be discovered by the naïve-theory theorist.

Stepping Out of the Box

Nonetheless, there are many exciting new discoveries—discoveries that do not resolve the dispute between theory-based and general-process based controversies—but that suggest instead that the field is moving in new directions, toward fundamental ideas about what categories and category development might be.

Perceptual Learning. New evidence shows that category learning systematically alters perceptual processing, creating dimensions and features and thus the very way the world is perceived and remembered. Quinn, Nelson and Snyder, and Gosselin and Schyn (chaps. 5, 1, and 4, respectively) all provide relevant data, much of it concerning human face perception. This is a domain in which one might expect to find hardwired competencies unaffected by experience. However, face-specific processing appears to be a consequence of experience with faces. As Nelson and Snyder argued, infants' early experiences with faces consist of seeing a very few individuals in many, many contexts. This learning environment appears to teach infants to attend to the configural properties that distinguish individuals. In contrast, infants' early experiences with common objects such as spoons, cups, and toys consist of encounters with many different instances of each category. Moreover, unlike people, instances of common categories are substitutable, that is, functionally equivalent. This learning environment appears to teach infants to attend less to the details of common objects and more to their overall similarities. This conclusion is supported both by evi-

dence from infants and from adult training experiments. In these studies, nonface stimuli were made to engage in face-perception processes by training participants to distinguish individuals (e.g., Tarr & Gauthier, 2000). Gosselin and Schyn's studies show that perceptual learning of this sort makes feature processing highly context- and task-specific. For example, adults process very different features when asked to judge the gender versus the emotion of a face.

Evidence in the literature suggests further that perceptual learning affects early as well as late stages of processing (see Goldstone & Barsalou, 1998, for a review). For example, effects of perceptual learning have been found in elementary perceptual tasks that precede decisions of category membership, including same-different judgment tasks (Goldstone, 1994) and part detection (e.g., Lin & Murphy, 1997). Other studies suggest that young children's attention to perceptual properties and perhaps even their parsing of object shape may change in fundamental ways as a product of learning object categories (see Smith, 2003; also Needham, Barrett, & Peterman, 2002).

Perceptual learning clearly complicates the traditional view (e.g., Keil, 1994) of perception (and perceptual similarity) as raw, unprocessed and "knowledge-less." If fundamental perceptual processes change, become tuned to specific tasks as a function of experience in those tasks, then those processes are themselves knowledge-laden. A visual system that processes faces and objects differently "knows" that these are different kinds. The task dependency of perceptual features challenges Murphy and Medin's (1985) idea that category-specific feature selection can only be explained by causal theories of category structure. In this way, findings about perceptual learning undermine a distinction between perception and knowledge.

Embodiment. The knowledge-laden and history-dependent nature of perceptual processes encourages new ideas about representation, and specifically the idea that these are transient emergent events that are close to the sensory surface (Barsalou, chap. 15, this volume). In this view, the internal language for thought, perception, and action are fundamentally the same and must be so if they are to mesh seamlessly in creating intentional acts. A now growing industry of results supports this view. These include findings that actions prime categories (Tucker & Ellis, 2001), that objects are recognized better when we put our bodies in positions consistent with our usual actions on those objects (Creem & Proffitt, 2001), and even that verb meanings are tied to eye movements (Richardson, Spivey, Barsalou, & McRae, 2003). Studies using fMRI and PET are also consistent with these ideas and show that the visual identification of artifacts engages cortical regions associated with the typical motor actions on those artifacts (Ishai, Ungerleider, Martin, & Haxby, 2000; Faillenot, Toni, Decety, Gregoire, & Jeannerod, 1997).

In developmental studies, Rakison (chap. 6, this volume) and Gershkoff-Stowe (chap. 8, this volume) also find evidence for the role of action and the context-dependency of the features that children process when categorizing. Both sets of studies show that young children are highly sensitive to the features that are relevant to acting on objects: to the wheels that allow things to be rolled, to the springs that make objects bounce. These studies suggest a way in which one might bring together the insights of naïve-theory theorists that some features matter more because of their causal status and insights about perceptual learning and the embodied nature of general memory processes. In addition, these findings along with those of Barsalou remind that categories are acquired and used in a physical world by physical beings who act in that world and perceive the consequences of their own actions.

Nested Time Scales. In one influential series of experiments, Quinn (chap. 5, this volume) showed that infants' categories of "cats" and "dogs" are the emergent product of the experiences in the task, and do not represent prior, represented knowledge about the two categories. Instead, transient memories formed in the real-time experiences in the task—the repeated presentations of seeing particular cats—create categories in *performance*. Further, Quinn has begun to trace how the contributions of the immediate input, the task, and long-term experiences change with development. As the long-term memory contribution becomes stronger, decisions become more stable across stimulus and context. In her chapter, Gershkoff-Stowe also traces how infants' actions on objects—their in-task experiences—combine with their previous experiences to create in-task categories.

These studies remind us that the behavior we see at any moment is the product of processes operating over nested times scales (Smith & Thelen, 2003). The processes relevant to a baby's pattern of looking in the Quinn experiments or to a child's pattern of sequential touching in Gershkoff-Stowe's experiments include the sensory input at the moment of the behavior. But they also include the immediately preceding events: the cats and dogs one has seen in the last several minutes in the experiment, the objects the infant has touched and made to bounce just previously. The longer history of the infant—the dogs and cats seen over a lifetime, the objects and actions experienced over a lifetime—also matter. The relevant time scale to understanding categorization will therefore be in seconds, in minutes, in days, and in years. A complete theory must not just consider processes at all these time scales, but also integrate them and understand they influence each other over time. This is the dynamic systems perspective heralded by Gershkoff-Stowe (see also Samuelson & Smith, 2000).

Johnson's research (chap. 2, this volume) also illustrates this integrative approach. Johnson begins with the assumption that infants possess a vast array of perceptual processes and that these tune and adjust themselves

through their interaction with the world, and in so doing create higher level processes, what one might call, for example, object perception. Johnson notes findings parallel to those with human infants in infant monkeys and the correlated changes in the underlying neurophysiology. These neurophysiological studies make concrete the idea of many components interacting at many levels of analysis, from contrast receptors, to neural pathways, to the correlations among features in the world. The multicausality of developmental process is itself a challenge to the traditional framework for studying categories that demands a single answer to “what categories really are?” Johnson concludes that “It seems likely that unity perception cannot be reduced to a limited set of principles” (p. 56). It seems likely that human categories cannot be so reduced either.

SO WHAT IS NEXT?

It is always difficult to understand major change when one is in the midst of it. Further hindsight is always better than foresight. Thus, one cannot confidently make predictions about where the field is going. But it does seem to be going. The change is perhaps driven by the stale nature of the debates within the traditional metaphor: features or theories? But it is also driven by new advances in understanding neural processes and their plasticity, in computational approaches to learning, in large-scale analyses of the statistical structure of the learning environment, and in fundamental processes such as perceptual learning. It may not be clear where research on categories is going, but it *is* moving, and that is good.

REFERENCES

- Andrews, G., & Halford, G. S. (2002). A cognitive complexity metric applied to cognitive development. *Cognitive Psychology*, *45*(2), 153–219.
- Armstrong, S. L., Gleitman, L. R., & Gleitman, H. (1983). What some concepts might not be. *Cognition*, *13*(3), 263–308.
- Barrett, S. E., Abdi, H., Murphy, G. L., & Gallagher, J. M. (1993). Theory-based correlations and their role in children’s concepts. *Child Development*, *64*(6), 1595–1616.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*(3), 211–227.
- Barsalou, L. W. (1993a). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. Collins & S. Gathercole (Eds.), *Theories of memory* (pp. 29–101). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Barsalou, L. W. (1993b). Challenging assumptions about concepts. *Cognitive Development*, *8*(2), 169–180.
- Barsalou, L. W. (2003). Situated simulation in the human conceptual system: Conceptual representation [Special issue]. *Language & Cognitive Processes*, *18*(5–6), 513–562.

- Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning & Verbal Behavior*, 11(6), 717-726.
- Bruner, J. (1986). *Actual minds, possible worlds*. Cambridge, MA: Harvard University Press.
- Chater, N., & Vitanyi, P. M. B. (2003). The generalized universal law of generalization. *Journal of Mathematical Psychology*, 47(3), 346-369.
- Creem, S. H., & Proffitt, D. R. (2001). Grasping objects by their handles: A necessary interaction between cognition and action. *Journal of Experimental Psychology: Human Perception & Performance*, 27(1), 218-228.
- Estes, W. K. (1986). Memory storage and retrieval processes in category learning. *Journal of Experimental Psychology: General*, 115(2), 155-174.
- Faillenot, I., Toni, I., Decety, J., Gregoire, M., & Jeannerod, M. (1997). Visual pathways for object-oriented action and object recognition: Functional anatomy with PET. *Cerebral Cortex*, 7(1), 77-85.
- Fischer, K. W., Kenny, S. L., & Pipp, S. L. (1990). How cognitive processes and environmental conditions organize discontinuities in the development of abstractions. In C. N. Alexander & E. J. Langer (Eds.), *Higher stages of human development: Perspectives on adult growth* (pp. 162-187). London: Oxford University Press.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. London: 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. (2003). Theory-based categorization in early childhood. In D. Rakison & L. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp. 330-359). London: Oxford University Press.
- Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, 123(2), 178-200.
- Goldstone, R. L., & Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition*, 65(2-3), 231-262.
- Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, 130(1), 116-139.
- Hampton, J. A. (2001). The role of similarity in natural categorization. In U. Hahn & M. Ramscar (Eds.), *Similarity and categorization* (pp. 13-28). London: Oxford University Press.
- Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, 12(Suppl. 2), 35-51.
- Jones, S. S., & Smith, L. B. (1993). The place of perception in children's concepts. *Cognitive Development*, 8(2), 113-139.
- Karmiloff-Smith, A. (1999). Taking development seriously. *Human Development*, 42(6), 325-327.
- Keil, F. C. (1991). The emergence of theoretical beliefs as constraints on concepts. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and cognition. The Jean Piaget Symposium Series* (pp. 237-256). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Keil, F. C. (1994). Explanation, association, and the acquisition of word meaning. In L. Gleitman & B. Landau (Eds.), *The acquisition of the lexicon* (pp. 169-196). Cambridge, MA: MIT Press.
- Keil, F. C. (2003). Categorisation, causation, and the limits of understanding: Conceptual representation [Special issue]. *Language & Cognitive Processes*, 18(5-6), 663-692.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211-240.

- Lin, E. L., & Murphy, G. L. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception & Performance*, 23(4), 1153–1169.
- Malt, B. C. (1994). Water is not H-sub-2O. *Cognitive Psychology*, 27(1), 41–70.
- McRae, K., Cree, G. S., Westmacott, R., & de Sa, V. R. (1999). Further evidence for feature correlations in semantic memory: Visual word recognition [Special issue]. *Canadian Journal of Experimental Psychology*, 53(4), 360–373.
- Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, 44(12), 1469–1481.
- Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 20(4), 904–919.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289–316.
- Needham, A., Barrett, T., & Peterman, K. (2002). A pick me up for infants' exploratory skills: Early simulated experiences reaching for objects using 'sticky' mittens enhances young infants' object exploration skills. *Infant Behavior & Development*, 25(3), 279–295.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 10(1), 104–114.
- Nosofsky, R. M., & Palmeri, T. J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, 5(3), 345–369.
- Piaget, J. (1970). *Genetic epistemology* (E. Duckworth, Trans.). New York: Columbia University Press.
- Richardson, D. C., Spivey, M. J., Barsalou, L. W., & McRae, K. (2003). Spatial representations activated during real-time comprehension of verbs. *Cognitive Science*, 27(5), 767–780.
- Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21–59). New York: Cambridge University Press.
- Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328–350.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573–605.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26(5), 521–562.
- Samuelson, L. K., & Smith, L. B. (2000). Grounding development in cognitive processes. *Child Development*, 71(1), 98–106.
- Sera, M. D., Elieff, C., Forbes, J., Burch, M. C., Rodriguez, W., & Dubois, D. P. (2002). When language affects cognition and when it does not: An analysis of grammatical gender and classification. *Journal of Experimental Psychology: General*, 131(3), 377–397.
- Slooman, S. A. (1997). Explanatory coherence and the induction of properties. *Thinking & Reasoning*, 3(2), 81–110.
- Smith, E. E. (1989). Concepts and induction. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 501–526). Cambridge, MA: MIT Press.
- Smith, E. E., & Medin, D. L. (1981). *Concepts and categories*. Cambridge, MA: Harvard University Press.
- Smith, L. B. (2003). Learning to recognize objects. *Psychological Science*, 14(3), 244–250.
- Smith, L. B., Colunga, E., & Yoshida, H. (2003). Making an ontology: Cross-linguistic evidence. In D. Rakison & L. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp. 275–302). London: Oxford University Press.
- Smith, L. B., & Thelen, E. (2003). Development as a dynamic system. *Trends in Cognitive Science*, 7, 343–348.
- Tarr, M. J., & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, 3(8), 764–769.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity and Bayesian inference. *Behavioral & Brain Sciences*, 24(4), 629–640.

- Tucker, M., & Ellis, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition*, 8(6), 769–800.
- Vigliocco, G., Vinson, D., Lewis, W., & Garrett, M. (2004). Representing the meanings of object and action words: The Featural Unity and Semantic Space (FUSS) hypothesis. *Cognitive Psychology*, 48, 422–488.
- Wohwill, J. F. (1967). The mystery of the pre-logical child. *Psychology Today*, 1(3), 24–34.
- Yoshida, H., & Smith, L. B. (2003a). Shifting ontological boundaries: How Japanese- and English-speaking children generalize names for animals and artifacts. *Developmental Science*, 6(1), 1–17.
- Yoshida, H., & Smith, L. B. (2003b). Response: Correlation, concepts and cross-linguistic differences. *Developmental Science*, 6(1), 30–34.
- Zaki, S. R., & Nosofsky, R. M. (2001). Exemplar accounts of blending and distinctiveness effects in perceptual old-new recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 27(4), 1022–1041.