In K. Lamberts & D. Shanks (eds.) Knowledge, concepts and categories Psychology Press, East Sussex, UK 1997

CHAPTER FIVE

Perceiving and Remembering: Category Stability, Variability and Development

Linda B. Smith and Larissa K. Samuelson

> . . . no two ideas are ever exactly the same, which is the proposition we started to prove. The proposition is more important theoretically than it at first sight appears (William James 1890:235).

Traditional theories of categorization concentrate on the stability of cognition, how it is that people perform the same cognitive act over and over despite varying local circumstances. For example, each time a person encounters a frog, whether in a pond or in a comic strip, the encountered object is understood to be the same kind of thing – a frog. Traditional theory explains this stability by positing stable mental representations. According to this view, the reason that people understand the frog in the pond and the frog in the comic strip to be the same sort of thing is that in both cases they access the same representation, the same concept of frog.

Figure 5.1 portrays an act of knowing in this view: a sensory event makes contact with the concept of frog and it is at the instant of contact – when perceiving activates permanent knowledge stores – that the individual may be said to "know" that it has seen a frog. Thus knowledge, in the form of relatively permanent representations, is distinct from knowing which consists of momentarily activated representations. Most research on categories has been concerned with the structures of these stable representations presumed to underlie moments of knowing.

The problem, of course, is that as empirical scientists we have only individual moments of knowing as our window onto knowledge. And

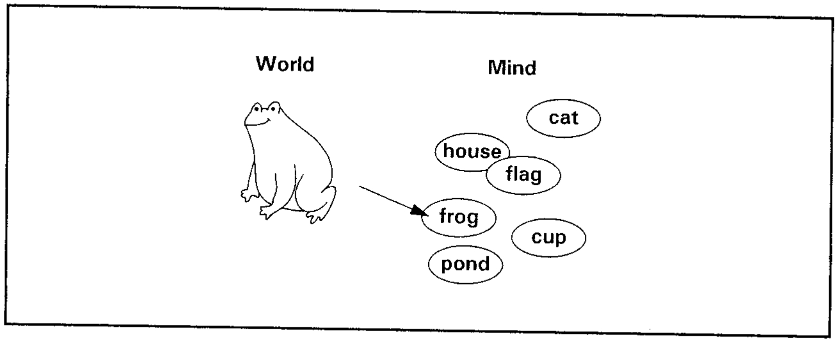

FIG. 5.1. The traditional view of a moment of knowing: an object makes contact with a represented concept.

these, unlike the presumed constant concepts that underlie them, are highly variable. For example, when we encounter two frogs on two different occasions, we do not think exactly the same thing. Barsalou (1987) demonstrated this empirically; he found that when people read the word *frog* in isolation, they do not think "eaten by humans". But when *frog* is read in the context of "French restaurant", "eaten by humans" does come to mind. What is "known" in a real moment of knowing depends on the context. As William James noted in the opening quote, this fact has profound theoretical consequences. We pursue these consequences in the following three parts of this chapter.

Part 1 briefly reviews the traditional search for constant concepts in the face of variable moments of knowing. It is not a history of success (see Barsalou 1989, 1987, Komatsu 1992, Goldstone et al. in press; and Jones & Smith 1993).

Part 2 explores how the processes, perceiving and remembering, that make individual acts of knowing also make categories. These results suggest a unified account of category stability and variability that does not make use of the idea of a represented concept.

Part 3 applies these ideas to category development and demonstrates the potential value of James' insight that "no two ideas are ever the same", for explaining developmental change.

PART 1: THEORIES ABOUT CONCEPTS

> What makes it convenient to use the mythological formulas is the whole organization of speech, which, as was remarked a while ago, was not made by psychologists, but by men who were as a rule only interested in the facts their mental states revealed. They spoke of their states as *ideas of this or of that thing*. What wonder, then, that the thought is most easily conceived under the law of the thing whose name it bears! (William James, 1890:236).

The traditional framework depicted in Figure 5.1 makes sense when one thinks about the stability of human categories in the face of a highly variable world. We regularly recognize quite diverse things as instances of the same kind: we know pond frogs and restaurant frogs and poison-dart frogs to all be frogs. A permanently represented concept has been seen by many to be necessary to explain this stability. For example, Keil (1994) wrote

> 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 (1994:169).

Thus, "repeated categorizations based on the same mental representation [are] the launching point for most psychological investigations of what concepts are".

The traditional framework depicted in Figure 5.1 also fits the understanding of categories found in 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". Psychological theories of categories reflect these ideas of fixed extensions and intensions. In psychology, the extension is the "repeated categorizations" that people make, i.e. the data to be explained. The intension is the hypothesized concept that determines the extension, the mental structure that causes people to categorize objects the way that they do.

These foundational ideas of stable categories and stable concepts, however, have led to little progress. Instead a steady succession of theories of concepts have been offered, rejected, resurrected, and rejected again.

From criterial properties to essences

In the 1960s, research on concepts concentrated on psychological 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 of such concepts (Murphy & Medin 1985, Fodor 1977). Secondly, 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 Smith & Medin 1981, Wittgenstein 1953, Katz 1972, Rosch & Mervis 1975). Thirdly, there were data that directly contradicted the idea of necessary and sufficient features. Specifically, if category membership is determined by necessary and sufficient properties then all members of the category are equally good. People, however, consistently and systematically judge some instances of a category to be better than others (Rosch 1973). A robin, for example, is a better bird than a blue jay.

Thus, in the 1970s, the field turned to probabilistic theories of concepts (Smith & Medin 1981). Concepts became lists of characteristic rather than defining features (or in exemplar models, lists of the features of each experienced instance). These probabilistic theories derived from and readily explained graded category judgements. However, probabilistic theories came to be rejected for the same reasons that the defining-feature theories were rejected. First, there is no principled basis for determining the features that comprise representations (see Murphy & Medin 1985). Secondly, no one has yet offered a successful accounting of any natural category. And, thirdly, people make category judgements that are not in accord with probabilistic feature representations. Specifically, people *believe* that categories are organized by defining features that are necessary and sufficient. For example, along with judging robins to be better birds than blue jays, people also believe that all birds possess some property possessed by all and only birds (Rips 1989, Keil 1994). Moreover, people will maintain that an object that does not look or act at all like a bird, but instead looks and acts like an insect with elongated, transparent wings and antennae, nonetheless "really is a bird". They do so if they are also told that this insect-like organism has birds for parents and has bird DNA (e.g. Rips 1989). In such judgements, being a bird seems to depend criterially on special properties that are causally related to the origins of the entity.

In light of these data, theories of concepts in the 1980s turned from feature-based (or similarity-based) representations to theory-like (or explanation-based) representations (Murphy & Medin 1985). Many researchers studying people's beliefs about what "really makes something what it is" proposed that relevant representations are more like naïve theories than lists of criterial or probabilistic features (Keil 1989). In this view, concepts are domain-specific belief systems about the causal relations among properties and among concepts. Thus, the

reason that an organism that looks and acts nothing like a bird might still "really be a bird" is because the organism possesses the essential property, the biological structure, that *causes* the outward manifestations characteristic of birds. These ideas resurrect the criterial-property concepts of the 1960s. However, by the more contemporary view, it is not that all instances of a category share an "essence" but that people believe that they do and use that belief when they decide whether something is or is not a member of a category.

Unfortunately, concepts as naïve theories have the same problems as previous theories about concepts. First, there is no consensus as to what naive theories are, how they are mentally represented, or what kinds of knowledge are included. There is no principled or psychological basis for deciding what counts as a "theory" and what does not (see Keil 1994). Secondly, there is no well-formulated account of any natural category within this framework. Thus, once again, there is no demonstration that this kind of theory can actually do the job of explaining human category judgements. Thirdly, and even amidst the vagueness of these accounts, naïve theories clearly do not explain all the data. For example, they offer no account of why robins are psychologically better birds than blue jays.

One contemporary response to this state of affairs is to propose that concepts are composed of multiple parts, with some parts organizing categorization in some tasks and other parts organizing category judgements in other tasks (see Keil 1994). Specifically, probabilistic feature associations have been proposed to be involved in the identification of instances as members of a category and theory-like knowledge has been proposed to underlie people's reasoning about categories. Figure 5.1 thus becomes Figure 5.2. However, recent evidence suggests that explaining human categorization may require a more radical departure from old ideas about concepts, certainly more radical than piecing together old theories into a "two-part" hybrid.

Variable categories

Malt (1994) contrasted the liquids people label *water*, how much of the psychologically essential property (H_2O) these potential instances of "water" had, and the typicality of each as an instance of the category "water". The principal result is that none of these category judgements aligned (see also Rips 1989, Rips & Collins 1993). Tea, for example, which is not called water at all, is judged to have a higher percentage of the essential property (H_2O) than ocean water which is uniformly called water and is judged to be highly typical of the category "water". Mineral water, in contrast, is called water, is judged to have a high percentage of H_2O, but is also judged to be an atypical instance of the

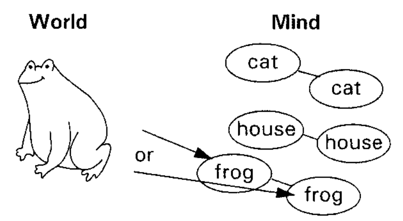

FIG. 5.2. Hybrid concepts: one part is about causes and essences and is accessed in some tasks and the other part is probabilistic feature lists (of experienced exemplars) and accessed in other tasks.

category "water". In brief, Malt's evidence suggests three distinct patterns of category judgements over the same domain of instances – one in naming tasks, one in judging the percentage of "pure water", and one in judging category typicality. Is it sensible to suppose that there are parts of a single concept devoted to each of these tasks?

Lakoff's (1987) analysis of the word "mother" poses the same problem for the idea of a single cohesive and coherent concept. By Lakoff's analysis, people sometimes use the word "mother" to refer to the person who gave birth; other times to the nurturing female adult; other times to the wife of the father. These various meanings are mutually contradictory as seen in such statements as, "She's not my real mother, all she did was give birth to me", and "I searched for my real mother because I wanted to know my genetic origins". Similar to Malt's findings about water, Lakoff's analysis of "mother" suggests contradiction, context-dependency, and multiplicity in the place of coherence, stability, and singularity.

Other evidence suggests shifts in category membership occur with quite small and seemingly non-meaningful shifts in contexts. For example, Rips (1989) found that people judged a three-inch diameter object to be a more likely member of the category "pizza" than the category "quarter" – a judgement that seems to contradict the greater intuition that three-inch objects are more similar to quarters than to typical pizzas. Rips justified subjects' judgements by pointing to the fixed size of quarters determined by the process of minting them. He posited that subjects' concepts were theory-like accounts of the origins of things. Recently, however, E. Smith and Sloman (1994) replicated Rips' study with one small change. Rips had told subjects about only one property

of the to-be-classified item, that it was three-inches in diameter. Smith & Sloman, in contrast, told subjects about two properties, that the to-be-classified object was three inches in size and silver-coloured. Everything else about the two studies was the same; but the results were just the opposite. In the Smith & Sloman study, people did not make use of their knowledge about the set sizes of quarters but instead judged category membership by overall similarity (that is, a three-inch silver object is judged to be a quarter and not a pizza).

Altogether Malt's, Lakoff's, and Smith & Sloman's results suggest that the fundamental assumption made by theories of concepts, what Keil called "shared representations" that make "repeated categorizations", is not quite right. Acts of categorization are not simply repeated; they vary. Different tasks and contexts seem to create different categories.

The plausibility of this last idea is demonstrated in Barsalou's (1983) elegant studies of *ad hoc* categories, Barsalou asked adults to determine "all the things on a desk that could be used to pound in a nail". Subjects readily formed novel and contextually-coherent categories. These *ad hoc* categories were not based on a previously represented concept; nonetheless, they show the same performance characteristics as more ordinary categories, that is, well-organized typicality judgements, graded category membership, and characteristic features. These findings present a strong challenge to the traditional view of knowing shown in Figure 5.1; they suggest that individual acts of categorization do not require an already represented concept.

Can a theory of concepts be saved?

Together these data revive James' earlier criticism of concepts as potentially "mythological entities". Two responses to the growing concern about the very idea of a concept argue that the notion of concept should be retained. One response asserts that categories, the data notwithstanding, are not variable. In this view, the existence of fixed taxonomies through which people understand their world is a self-evident truth; that psychologically, individual entities either are or are not water, either are or are not mothers. This position can be held if one assumes that some of the tasks that psychologists have considered to be category-judgement tasks are not really category-judgement tasks at all. Armstrong et al. (1983), Gelman & Coley (1991), Soja et al. (1991), and Rips (1989) have all argued that careful introspections by subjects about category membership reveal the fixed taxonomy through which people represent the world. Other judgements such as object recognition, typicality judgements, naming, and similarity judgements, in contrast, are not really category judgements in that they do not rely on the represented taxonomy. The problem with this argument is that it dismisses as irrelevant rather than explains the findings of category variability.

A second response arguing for the idea of a concept has been offered by Medin & Ortony (1989). They argue that stable concepts might be inputs to the processes that create variable categories. Medin & Ortony wrote

> we think care has to be taken not to equate instability in outputs or behaviours with underlying or internal instability. Might it be that underlying concepts are in fact stable (whatever that might mean) and that the apparent instability is an artifact of the processes that operate on these stable representations? (1989:191).

This view is pictured in Figure 5.3. We have a single coherent concept of, for example, frog that we access whenever we encounter or think about frog – when we recognize a frog, when we introspect about the essential properties of frogs, when we make typicality judgements, and when we reason about frogs in ponds versus frogs in restaurants. In this view, real-time processes of perceiving, remembering and responding to specific task demands make individual acts of categorization different. But underneath all the variability of individual acts is the very same concept.

This proposal explicitly recognizes two main facts about human categorization that need to be explained: the stability across individual acts of categorization and the uniqueness of individual acts. Before we consider further the viability of this idea, there is a third fact about human categories that needs to be explained.

Category development

Acts of categorization do not just vary from task to task, they also vary over developmental time. Eighteen-month-old children, but not adults,

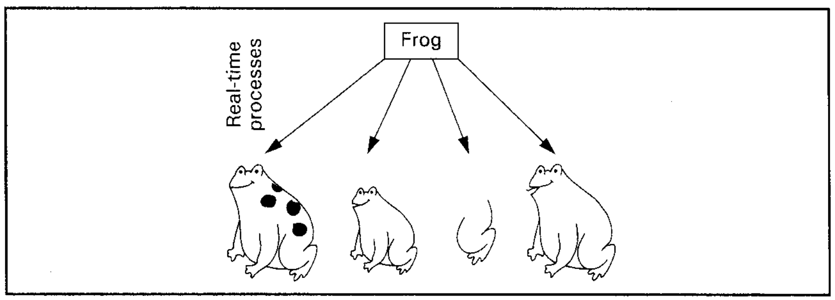

FIG. 5.3. Medin & Ortony's idea of a constant concept that gives rise to variable performance.

sometimes classify goats with dogs by calling a goat "doggie" (Mervis 1987). Pre-school children, but not older children and adults, maintain that the category "island" includes landmasses with beaches and palm trees that are not surrounded by water (Keil 1989). Three-year-olds, but not adults, call objects "comb" that look like combs even if they are made from materials that cannot possibly untangle hair (Landau et al. submitted).

Given the theoretical framework of concepts, the principal question developmentalists have asked about these kinds of results is whether immature intensive definitions are fundamentally different from more mature ones. In the 1960s, research focused specifically on whether pre-school children's categories were based on illogical complexive definitions in contrast to the presumed criterial property structure of mature concepts (Inhelder & Piaget 1964). In the 1980s and 1990s, the question shifted to whether young children's concepts are organized by probabilistic clusters of perceptual properties in contrast to the presumed theory-based concepts of adults (Keil 1989).

Just as there is no strong evidence for any single kind of concept in the adult literature, there is no strong evidence for a developmental shift from one particular kind of concept to another. Indeed, although children's category judgements are different from adults, they appear to be like adults' category judgements in being contextually and task determined. One compelling set of findings on the contextual nature of young children's categories are those of Imai et al. (1994). In one task, they presented three-year-old children with an exemplar object, for example, a picture of a birthday cake and then asked the children which of three test objects "go with" the exemplar: (a) a same shaped object (a wide-brimmed hat); (b) a taxonomically-related object (a pie); or (c) a thematically-related object (a birthday gift). Children predominantly chose the thematically related object. In a second task, children were asked to generalize a novel "dinosaur talk" name. For example, they were told that in dinosaur talk, the birthday cake was called "dax". The children were then asked which of the three test objects was also called "dax". The children primarily chose the same shaped object, the hat. In this same task, older children and adults extended the novel name to the taxonomically related object. Although Imai et al. did not elicit taxonomic-category judgements from these children, evidence from other researchers strongly suggests that if the children had been asked to feed the dinosaur, they would have chosen the cake and the pie (Bauer & Fivush 1992). Altogether, this evidence suggests that children categorize objects differently when making "goes with" judgements, when naming, and when making functional judgements (see also, Landau et al., submitted; Smith et al., 1996). In brief, children's categories, like those of adults, are contextually variable. But critically, they are also contextually variable in different ways than those of adults.

In light of this evidence, reconsider Medin & Ortony's suggestion, illustrated in Figure 5.3, that stable concepts serve as input to variable processes. How would such a model explain concept variability both across tasks and across development? The theoretical problem is illustrated in Figure 5.4. The figure depicts a represented concept at two points in development and the variable performances that emerge from these two representations. In this picture, two distinct kinds of variability need to be explained – change in concepts with development and changes in individual acts of categorization across contexts and tasks. In Part 2, we lay the groundwork for an alternate solution, one in which the very same processes that make variability across individual acts of categorization make variability across developmental time.

Summary

Theories of concepts have concentrated on the stability of categories – on the fact that people treat quite diverse entities as equivalent and that they do so in globally similar ways across contexts and tasks. However, the evidence suggests that on closer inspection categories are variable as well as stable. Further, people appear able to create categories on the spot. Category variability and in task category creation are facts not well explained by the idea of a concept. In the next section, we seek insights into these phenomena from what is known about processes of perceiving and remembering.

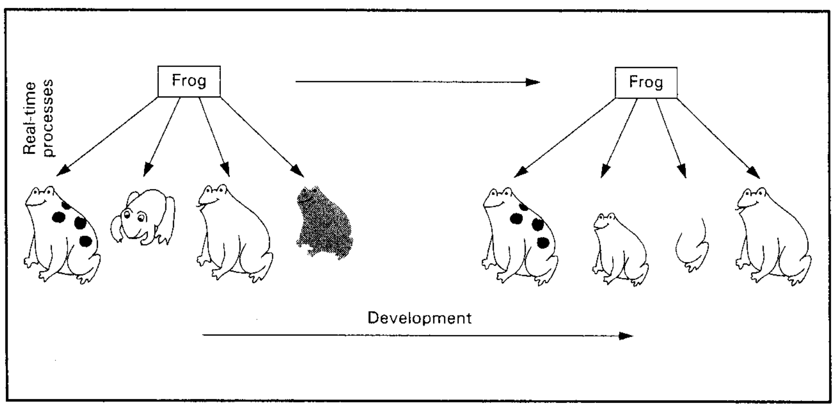

FIG. 5.4. Cartoon illustrating concepts that change with development and also that lead to contextually-variable performances.

PART 2: PERCEIVING AND REMEMBERING

> For an identical sensation to recur it would have to occur the second time in an *unmodified brain*. But as this, strictly speaking, is a physiological impossibility, so is an unmodified feeling an impossibility (William James 1890:232–3).

> . . . brains use *processes that change themselves* – and this means we cannot separate such processes from the products they produce. In particular, brains make memories, which change the ways we'll subsequently think. *The principal activities of brains are making changes in themselves*. (Minsky 1986:288).

Figure 5.1 with which we began this chapter presents one picture of knowing: seeing a frog activates a stored knowledge representation that provides the meaning of the sensory event. This picture foregrounds the role of stored knowledge in knowing. Backgrounded are the processes – perceiving and remembering – that realize that knowledge in a particular moment of knowing. In this part, we turn to the considerable evidence on processes of perceiving and remembering for insights into the stability and variability of categories. This extensive literature points to three fundamental truths about cognitive processes:

(1) they depend broadly on the immediate input and its larger context;

(2) they are temporally extended with real rise times and decay times such that activity at any moment depends on and emerges out of preceding activity;

(3) they change as a direct consequence of their own activity.

In this section we review the evidence. We then show how the stability and variability of human categories arise naturally from these facts.

The immediate context

The first fact about perceiving and remembering is contextual dependency. This context dependency is at direct odds with classic psychophysical approaches to perception (see Marks 1993, Schyns & Rodet, in press). The classic work sought to map invariant components of the physical world that gave rise to invariant percepts – the wavelength that specified a colour, the voice onset times that specified a phoneme, the features that specified a letter. We now know that perception is much messier than this; there may be no finite set of primitives out of which perceptions are built, no 1:1 map from specific inputs to specific percepts. Instead, perception is relative, and contextually determined: the reflected wavelengths seen as red when they emanate from firetrucks are the not the same as the ones seen as red when they emanate from hair (Halff et al. 1976); the voice onset times that signal a certain consonant are not fixed but depend on other aspects of the stimulus including the rate of talking (Kelly & Martin 1994); the relevant features for letter perception shift radically with the font (Sanocki 1991).

Figure 5.5 illustrates several more classic examples of the role of context in perception. Panel a shows how the perceived size of an object depends on surrounding objects: an object looks smaller than it really is when surrounded by objects that are much larger yet looks larger than it really is when surrounded by objects just slightly larger than itself. Panels b and c show how the perceived similarity of two objects depends on other perceptually present objects. In Panel b, the perceived similarity of objects 1 and 2 is low but the perceived similarity of these same two objects in Panel c is high (for more relevant data on the context-dependency of perceived similarity see Goldstone et al. 1991). Finally, Panel d shows how the addition of a constant line, a small change in context, radically transforms shape and perceived similarity (Palmer 1989).

The importance of these textbook facts about the context-dependent nature of perception should not be underestimated. They mean that the psychological object, the object to be categorized, is not itself a fixed

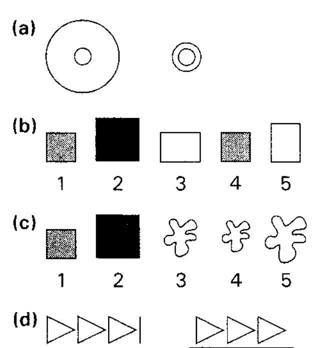

FIG. 5.5. Illustrations of context effects in perception: a – perceived size of centre circle depends on size of surrounding circle; b and c – perceived similarity of objects 1 and 2 depend on other objects in the comparison set; and d – the perceived shape and orientation of triangles depends of the perceived frame.

entity with one objectively correct description. Rather, psychological objects, like the categories in which we place them, depend in a chameleon-like way on the surrounds.

The contextual malleability of general cognitive processes is also seen in many memory phenomena. Light & Carter-Sobel's (1970) classic demonstration of encoding specificity provides one good example. They showed that the word "jam" encountered in the context of traffic does not lead to the same memory as the word "jam" encountered in the context of strawberry. Further, the to-be-remembered word, jam, is better recognized by subjects in the context that matches the original learning (the word "traffic") than one that is different. The importance of context goes far beyond this paradigm (see Tulving & Thomson 1973). Evidence from a variety of memory tasks indicates that what is remembered depends critically on a holistic match between the quite general context of the original event and the context of the moment. Thus, Godden & Baddeley (1980) found that scuba divers who learned lists of words underwater remembered them better when tested underwater than on land. Similarly, Butler & Rovee-Collier (1989) found that babies who had learned to kick to make a mobile bounce in a crib with a particularly patterned crib sheet and bumper, remembered days later what they learned when the crib sheet and bumper were the same but not when they were different. Other evidence shows that the particular room, the particular voice of a speaker, and even the mood of the subject matter in what is remembered and what is recalled (S. Smith 1986, Palmeri et al. 1993, Eich 1985). In sum, what we remember depends broadly on the moment of learning and the moment of retrieval.

These facts about perceiving and remembering have profound implications for theories of categories. Objects can not map stably into a set taxonomy because the objects of mental life are not themselves stable entities. These facts also mean that mental events are naturally adapted to context – thoughts about frogs in restaurants differ appropriately from thoughts about frogs in ponds because the perceived object and the remembered knowledge are made in and from the context of the moment.

Continuity with the just-previous past

The second fact about perceiving and remembering is that these processes are extended in time. Because the information-bearing events that comprise perceiving and remembering take time and endure, mental activity at any point in time will be a mixed result of immediate input and just-past activity. This is readily seen in tongue twisters. The difficulty in saying "Peter Piper picked a peck of pickled peppers" lies not in saying "pickled" but in saying "pickled" after "picked". The journals and textbooks of cognitive psychology are filled with many more examples.

One such example is the ubiquitous phenomena of "priming": within a narrow time frame, the perception of a prior word (or object) facilitates the perception of subsequent words (or objects) that are similar in some way. Thus perceiving the word "doctor" facilitates perceiving "nurse". Phenomena of "priming" are widespread in lexical processes, picture perception, object recognition, and motor behaviour (see for example, Klatzky 1980, Rosenbaum 1991, Harris & Coltheart 1986).

All theories of priming, in one way or another, posit that the internal activity that gives rise to the recognition of the first item (e.g. doctor) persists and thus facilitates processing of the second. The basic idea is that the pattern of activity associated with the first item overlaps in kind with the pattern underlying the perception of the second item and thus puts the second in a state of partial activation and readiness. Priming shows without a doubt that the thoughts we have at one moment grow out of those just before.

The temporal reality of cognitive processes, the shaping of the present by the just-previous past is also everywhere evident in perception. One class of relevant phenomena are adaptation effects: the repeated presentation of an event alters the perception of subsequent events. Such adaptation effects are seen in what one might think of as the primitive sensations of colour, loudness and pitch (e.g. Anstis & Saida 1985, Marks 1993), but also in the more complicated perceptions of musical chords (e.g. Zatorre & Halpern 1979), shapes (e.g. Halpern & Warm 1984), speech sounds (e.g. Remez et al. 1980) and faces (O'Leary & McMahon 1991). For example, prolonged staring at curved lines causes physically straight lines to be perceived as curved (Gibson 1933). The repeated presentation of a sound can cause subsequent sounds to be perceptually assimilated to it or to be perceptually pushed away from the adapting event (Marks 1993). The repeated presentation of one stimulus event can even shift what might seem to be preset category boundaries. For example, Remez (1979) shifted the boundary between whether sounds were perceived as a vowel or a buzz by the repeated presentation of an |a| sound.

These ever-present adaptation and priming effects, like the pervasive context effects, mean that mental states are extremely unlikely ever to be repeated. They mean, as James put it in the opening quote of this chapter, that we never have the same idea twice. These ideas also mean that there is a pull for coherence from one thought to the next one, for the meaning of an event to depend on its place in a stream of events. If we think first about eating and then about frogs we will think differently than if we think first about ponds and then about frogs.

A history of individual acts of knowing

The third fact about processes of perceiving and remembering is that they change themselves. An act of perceiving or remembering causes not only transient changes but also longer lasting, near permanent, changes. We know long-lasting changes must happen or we would have no memories of the individual events of our own lives and no connectedness with our own past. Empirical evidence suggests further that the power of a single processing event to alter subsequent knowing can be quite remarkable. We mention two such examples.

One example is Jacoby et al.'s (1989) ability to make people famous overnight. They had subjects read a list of names that included all non-famous people, names such as Samuel Weisdorf. Twenty-four hours later they gave subjects a list of famous and non-famous names and asked subjects to pick out the famous people. Subjects picked out Samuel Weisdorf along with Minnie Pearl and Christopher Wren. Having read the name once was sufficient to create a lasting degree of familiarity – one sufficient for a categorization of the name as "famous".

A second example is Perris et al.'s (1990) equally dramatic demonstration of toddlers' memory of a single experimental session that occurred in their infancy. The original experimental event was designed to test infants' use of visual cues to control reaching. To do this Perris et al. (1990) taught six-month-old children to reach in the dark for different-sized objects. The different sizes were signalled by different sounds (e.g. bells for big objects, squeaks for little ones). One to two years after the original experiment, Perris et al. brought these children back to the laboratory. At this point, the children were between 18 and 30 months of age. At this test session, the lights were simply turned off, the sounds played and the children's behaviour was observed. Perris et al. found that the children who had been in the experiment as babies reached in the dark for the sounding objects; control children who had not participated in the infant study did not. Thus, the one-time experience at six months permanently changed these children, altering the likelihood of behaviours one and two years later.

There are many more such demonstrations of long-lasting facilitatory effects in the literature – of the benefits of a single prior processing experiences (with units as small as single words) that have effects days, weeks, years later (e.g. Jacoby 1983, Salasoo et al. 1985, Rovee-Collier et al. 1985, Brooks 1987). These results indicate that each act of perceiving and remembering changes us.

Critically, the accrual of these long-term changes provides a source of stability in a continually changing system. If there are statistical regularities, patterns, in our experiences that recur over and over again, then as each moment of knowing is laid on the preceding moments, weak tendencies to behave and to think in certain ways will become strong tendencies – sometimes so strong that they will not be easily perturbed and thus might seem fixed.

There is considerable evidence that people are ready learners of statistical regularities. Indeed, people appear to learn whatever sorts of regularities are presented to them (Kelly & Martin 1994, Hasher & Zacks 1984, Coren & Porac 1977, Saegert et al. 1973, Shapiro 1969, Ashby et al. 1993, Lewicki et al. 1989, Reber et al. 1980). Examples include the pervasive effects of word frequency on word recognition (Harris & Coltheart 1986), typicality judgements which often reflect the most commonly-experienced instances (e.g. robin) and the strong effect of the frequency with which a face or brandname has been experienced on the likeability of that face or brandname (Zajonc 1968). Other evidence includes adults' remarkable sensitivity to the frequencies of events in the world – from the lethality of different events (Lichtenstein et al. 1968) to the frequency of fast food restaurants (Shedler et al. 1985; see Kelly & Martin 1994, for a review).

Still other evidence indicates that even very young children are highly sensitive to statistical regularities and, moreover, that these regularities play a demonstrable role in their category boundaries. For example, Sera et al. (1988) showed two-year-olds series of objects varying widely in size – for example, sneakers that varied from doll size to US men's size 18, buttons that varied from dots to platter size, and plates that varied from button size to several feet in diameter. In each case the children were asked individually about each object whether it was big or little. For each category, sneakers, buttons, and plates, the children imposed a sharp boundary between the sizes designated big and those designated little. For each category, that boundary fell at the (likely) most commonly experienced size by the child for that category – their own shoe size for sneakers, the size of a shirt button for buttons, and the 12 inch standard dinner plate size for plates. Apparently, children's everyday experiences with specific objects – putting on their own shoes, having their shirts buttoned, sitting at the dinner table – create in aggregate quite good knowledge about the specific sizes of specific kinds of things.

Other evidence suggests that regularities, correlations among multiple properties, may shape category judgements in unexpected ways. For example, Sera et al. (1994) asked speakers of Spanish and English to "cast" an animated film in which everyday objects came to life. The subjects' specific task was to decide whether the voices for these animated objects should be male or female. Spanish and English speakers provided quite different castings. For example, the Spanish speakers classified arrows and wheels along with ballerinas and queens as female voices but classified ice cream and shoes along with kings and giants as male voices. In contrast, English speakers cast arrows and wheels as male voices, ice cream as a female voice and shoes as possibly either. Critically, the Spanish speakers' judgements were predictable from the grammatical gender of the lexical item. Apparently, the association of shoes and ice cream with kings and men via the determiner "*el*" (as opposed to "*la*") makes shoes and ice cream more manly for Spanish than English speakers.

Widespread sensitivity to all forms of regularities and patterns seems likely to play a crucial role in categorization generally. As Kelly & Martin (1994:107) wrote, "the world is awash with stuff best described as 'tendencies', 'maybes', 'estimates', and 'generally speakings'". No individual regularity may be enough to explain the stability and context sensitivity of categories, but the combination of the many imperfect relationships in the world may. Consistent with this idea, Lakoff (1987) suggested that grammatical gender categories stably emerge and are productive because they are made of imperfect mixtures of imperfect cues – biological gender, cultural associations, the perceptual properties of objects, phonological properties of the word. Similarly, Kelly (1994) argued that parsing speech into word units may be dependent on a complex web of prosodic, phonological, and morphophonemic cues.

Statistics, however, may not be everything. The history, or specific order of events in a learner's past may be a critical determiner of what is learned. One provocative demonstration of this idea is Schyns & Rodet's (in press) study of how category learning may create new perceptual features. In their study, adults learned about two different types of cells from a Martian creature; Cell type A or Cell type AB, illustrated in Figure 5.6. Cell type A was defined in terms of a single feature, illustrated by the arrow in the figure. Instances of cell type A differed in other components but all possessed this critical feature. Cell type AB was defined by the more complex feature indicated by an arrow in the Figure. The critical manipulation was which specific category, type A or type AB, was learned first. The critical question was whether subjects defined cell type AB by a single feature that spatially conjoined the two parts or by two potentially separable features. After familiarization with the two categories, subjects were presented with test cells that spatially separated the two possible parts of the complex feature as illustrated in Figure 5.6. Subjects who learned cell type A first said the test cells were instances of type AB because they contained the two critical parts. Subjects who learned cell type AB first said test cells were instances of type A because they contained the A feature but not the complex B feature. Apparently learning A first had enabled

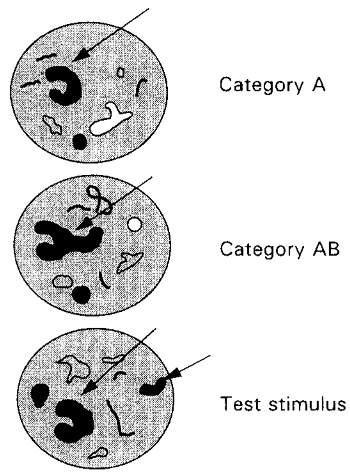

FIG. 5.6. Redrawn illustrations of stimuli used by Schyns & Rodet (in press). Arrows point to the critical features.

the subjects to parse the complex AB feature into two parts – the feature defining A and a new feature.

What is important about these results is that they show that what one knows depends not just on the total aggregate of learning experiences but their order. (See Regier 1995, for a computational model which provides a similar demonstration in the domain of spatial terms.)

How these facts may make categories

The extensive evidence on the contextual nature of perceiving and memory, the temporal groundedness of cognition in preceding activity, and its sensitivity to the history of its own activity provide a solution to the problem of how categories are both globally stable and locally variable, how novel categories may be created on-line, and how categories change over the life of an individual. Consider again the individual act of categorization that occurs upon seeing a frog and recognizing it as such. One's thoughts upon seeing the frog will be the combination of the immediate input in its full complexity, one's just preceding cognitive activity, and one's lifetime history of activity. The compression of all these sources of information in a single act means that what we know at a moment is an adaptive mix of the same stable regularities that also form other moments of knowing and the idiosyncracies of this moment.

We can see these ideas at work in three recent studies of how categories adapt themselves on-line. In one study, Goldstone (1995) asked subjects to judge the hue of objects by adjusting the colour of one object (the target) until it matched precisely another (the standard). The individual objects were letters and numbers and they were presented in a random order to the subjects to be judged. Unbeknown to the subjects, Goldstone had arranged for the colours and objects to be correlated across trials as shown in Figure 5.7. Specifically, the letters tended to be redder than the numbers. This fact strongly influenced subjects' judgements. Specifically, they judged presented letters (e.g. the "L" in Figure 5.7) to be redder than numbers of the exact same hue (e.g. the "8" in Figure 5.7). Apparently, subjects' lifetime history of experience with letters and numbers caused same category members to influence each other in the here and now. Long-term category knowledge combined with the transient effects of seeing redder letters than numbers and with the sensory information presented by the single to-be-judged object. Processes operating over different timescales combined in a single moment of knowing to make an individual letter look a particular degree of red.

The semantic congruity effect provides a second example of how knowing in a moment is made in the combination of long-term changes, transient in-task effects, and the immediate input. The semantic congruity effect refers to the finding that in comparative judgements, people are faster when comparing objects on a quantitative dimension when the direction of comparison is congruent with the location of the stimuli on the continuum. For example, when asked to make judgements about the size of animals, subjects are faster to choose the larger of two relatively large animals (e.g. elephant versus hippopotamus) than to choose the larger of two relatively small animals (e.g. hamster versus gerbil). Conversely, subjects are faster to choose the smaller of two small animals than to choose the smaller of two relatively large animals (Banks & Flora 1977). This general and robust effect clearly depends

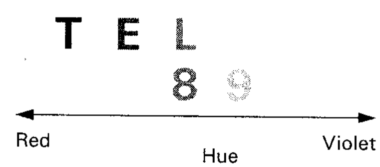

FIG. 5.7. Illustration of the correlations instantiated in Goldstone's (1995) experiment on colour perception.

on people's long-term and stable knowledge about the sizes of things. Banks & Flora originally suggested that people represent "elephant" as "very big" and therefore can answer questions about bigness directly. By this account, judging that an elephant is small (in comparison to say a whale) is difficult because it requires one to override the represented attribute "very big".

This account, however, cannot be the whole story. Our long-term knowledge about the sizes of things is not all that matters in the semantic congruity effect. For example, Cech & Shoben (1985) showed that the direction of the semantic congruity for a pair such as "rabbit–beaver" changes depending on the other pairs being judged in the task. In the context of other pairs of animals varying widely in size (from elephants to mice), "smaller than" judgements were faster than "larger than" judgements for the pair "rabbit–beaver". However, when this same pair was judged in an experiment in which rabbits and beavers were the biggest animals judged, the reverse semantic congruity was found: subjects were faster at judging which member of the rabbit–beaver pair was larger rather than smaller. Subsequent experiments have shown that the direction of semantic congruity for a given pair will shift in the course of an experiment as the sizes of the objects judged prior to the pair shift in one direction or the other (Cech et al. 1990). Thus, the semantic congruity effect is not dependent solely on the absolute value of a given item nor our long-term knowledge of the sizes of things. Rather, how fast one answers the question "Is this bigger than that?" depends on long-term knowledge, the preceding items just judged, and the immediate question asked.

The creation of transient "concepts" that meld the information from immediate input, from just-previous activity, and from a lifetime of activity seems fundamental to intelligence. One final example that makes this point is Sanocki's research on people's ease of recognizing letters in quite different fonts (1991, 1992; see also McGraw & Rehling 1994). The traditional approach to letter recognition (as we saw in the traditional approach to category recognition generally) is to try to specify the features that specify a particular letter – the features for example that enable one to recognize all the various letter "ys" in Figure 5.8. Sanocki's results, however, suggest a single set of represented and abstract features are not what enables us to recognize the letters of distinct fonts. The evidence against such an idea is that people are faster to recognize letters in familiar than in novel fonts and are faster to recognize letters consistent with the fonts of just previously seen or surrounding letters – even when the fonts are very well known. Altogether, the results suggest that people adjust their definition of features on-line to fit the font they are reading; for example, a specific

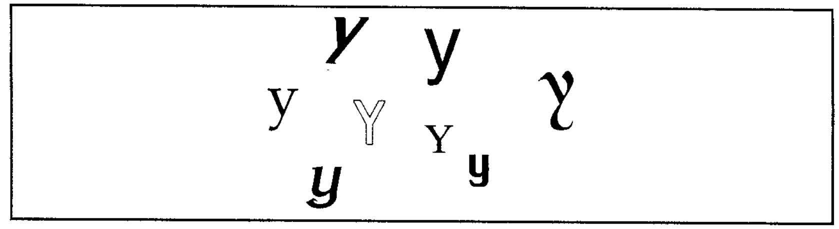

FIG. 5.8. The letter "y" in various fonts.

"y" that is difficult in one context is easy in another. Again, what we perceive when presented with a particular letter depends all at once on the immediate character of that letter, the character of the just-previously-perceived letters, and one's long-term experience of perceiving particular fonts.

We believe that these three examples – the influence of category knowledge on colour perception, the semantic congruity effect, the perception of letter categories – provide the key to understanding human categorization more generally. They suggest that categories exist only as the products of mental activity – in individual mental events with real-time durations that are themselves the product of their own lifetime of activity, the just-previous activity, and the immediate input. Categories that are this – the on-line product of complex processes of perceiving and remembering – will be dynamically stable, adaptive and, given an idiosyncratic mix of past and present, inventive. These ideas offer an alternative to the dualistic treatment of category stability and variability in Figure 5.3 and the dualistic treatment of real time processes and development in Figure 5.4.

We depict the new framework schematically in Figure 5.9. The activity of many heterogeneous and interacting subsystems that comprise a moment of knowing is represented by $*t$. The material causes of the activity at a single moment of knowing are the immediate input, the just-previous activity, and the nature of the cognitive system itself. The immediate input to the system at a particular moment in time is represented by I_t. The multiple processes of perceiving and remembering are indicated by arrows between the input and the individual moment of knowing and between one moment of knowing and the next. Importantly, since the activity at $*t$ is in part determined by the activity at $*t - 1$, it is also partly determined by the activity at $*t - 2$, $*t - 3 \ldots *t - n$. Each moment of knowing thus brings with it the history of its own past activity. Further, since each act of knowing permanently changes processes of perceiving and remembering, the accrued activity changes the cognitive system itself. It will not be the

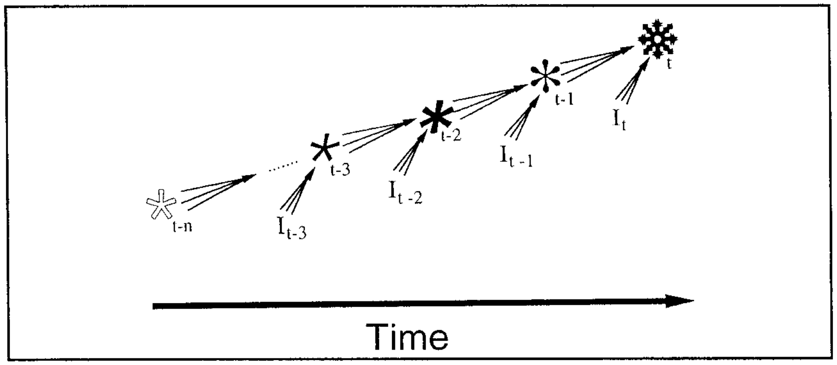

FIG. 5.9. Illustrations of how individual moments of knowing (*) combine immediate input (I), just-previous activity, and the history of activity.

same at t as it was at $t_{(n-1)}$. Thus real time and developmental time are unified; the very processes that make categories vary from moment to moment also make development.

PART 3: LEARNING NAMES FOR THINGS

> . . . our brain changes, and that, like aurora borealis, its whole internal equilibrium shifts with every pulse of change. The precise nature of the shifting at a given moment is a product of many factors. . . . But just as one of them certainly is the influence of outward objects on the sense-organs during the moment, so is another certainly the very special susceptibility in which the organ has been left at that moment by all it has gone through in the past. (William James, 1890:234).

In this section, we show how the ideas illustrated in Figure 5.9 both fit and are supported by data on children's category formation. We focus specifically on young children's initial generalization of a novel noun to new instances. Told that one object is, for example, a "dax", what other objects do children take to also be a "dax"? The now considerable evidence on this categorization task is critical to theories of categorization generally for three reasons. First, very young children in this task, like the subjects in Barsalou's *ad hoc* category experiments, appear to form coherently structured categories *on line*. Secondly, these categories seemingly created on the spot from hearing one object named once are often right, smart from an adult point of view. Thirdly, these phenomena seem central to the origins of lexical categories such as bird, frog, water – the ones that dominate the adult literature on concepts.

The critical results derive from studies employing an artificial word learning task. For example, in one such study, Landau et al. (1988) presented two- and three-year-old children with a small, blue, wooden inverted U-shaped object. They told the children that this exemplar object "is a dax". They then asked the children what other objects were also "a dax".

Figure 5.10 depicts the exemplar and the test objects. Given these stimuli, the children systematically generalized the name only to test objects that were the same shape as the exemplar – as if they already knew that objects of this shape (but necessarily a particular colour or material) were the same kind of thing. This result has now been replicated many times in many laboratories. Importantly, however, in these tasks children do not just form categories organized by shape. The nature of the categories they form depends on the immediate context, the just-previous events, and the child's history of naming things.

Contextual factors

Young children form well-organized categories specifically in word learning tasks. They do not do so generally in other kinds of categorization tasks. For example, given the objects in Figure 5.10, two- and three-year-olds do not categorize by shape when asked to make similarity judgements (Landau et al. 1988). Instead, they form categories based on holistic similarity or changing criteria. In a further

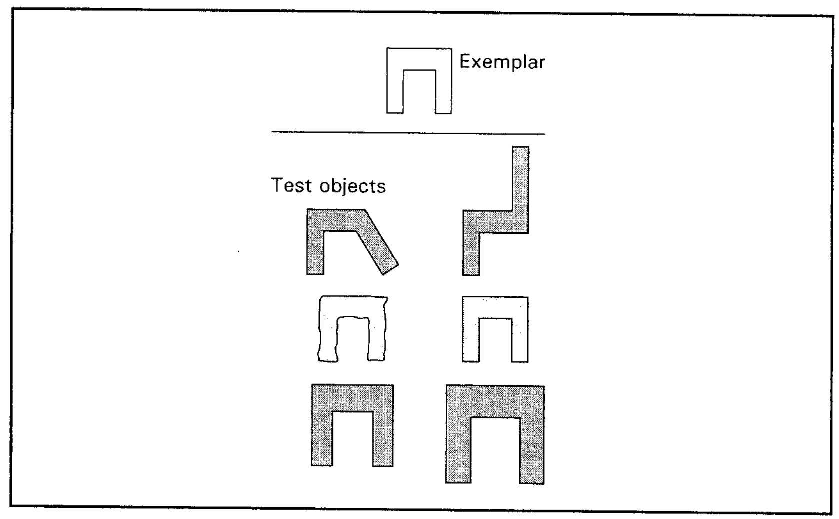

FIG. 5.10. Illustration of stimuli used by Smith et al. (1988): three-dimensional objects that varied in shape, size, and material (wood, sponge, and wire mesh).

study, Landau et al. (submitted) showed that children also classify differently when asked the names of things as against when asked to make judgements about function. For example, in one experiment, they used stimuli shaped like those in Figure 5.10, except now the exemplar and some of the test objects were made of sponge. The exemplar was named ("this is a dax") and then the experimenter spilled water on the table and wiped it up with the spongy exemplar. The children systematically generalized that name to same-shaped objects but when asked about function, these same children selected objects made of sponge when asked to wipe up water. The finding that systematic categorization by shape given stimuli like those in Figure 5.10 is specific to the task of naming has been replicated a number of times in several laboratories (e.g. Imai et al. 1994, Soja 1992). These findings tell us that stimulus properties alone do not compel the categories children form.

The stimulus properties, however, do matter. Children spontaneously form lexical categories by shape only when presented stimuli like those in Figure 5.10; changes in several kinds of properties change the categories children form. If the named objects have eyes, for example, children form lexical categories organized by shape and texture (Jones et al. 1991). If the objects are made of non-rigid substances (e.g. shaving cream with gravel in it), children form lexical categories organized by colour and texture (Soja et al. 1991).

What process creates these context effects on children's category formation in the task of generalizing a newly-learned word? Smith (1995, see also Jones & Smith 1993) suggest that it is a mix of here-and-now and learned forces on selective attention. The specific proposal is that memories of novel words are formed which include the properties of the object (and context) that children attend to at the time of hearing the novel word. The properties that children attend to will depend on the intrinsic salience of specific object properties and past attentional learning. This last point is a critical part of the proposal. One hundred years of research on attention suggests that the regular association of contextual cues with attention to some property leads to automatically increased attention to that property in that context. The language learning task presents the child with many statistical regularities which can guide attention and thus also guide language learning itself. There are perhaps associative relations between the syntactic frame "this is a _____" and specific object properties and also associative relations among object properties. In this view, the categories children form upon hearing a novel object named seem right from an adult point of view because they reflect, at least in part, statistical regularities in how words map to objects.

In sum, children's systematic generalizations of novel words used to label a novel object suggest that when children hear a novel object name, they form a category on the spot. Moreover the nature of the formed category appears to be influenced by a number of contextual factors – the task, the objects, and the linguistic context in which the novel word is embedded.

Continuity with the just-previous past

A growing literature indicates that young children's interpretation of novel words are also strongly influenced by the events occurring just prior to hearing the novel object named (see e.g. Baldwin 1991, Baldwin & Markman 1989, Tomasello & Kruger 1992, Olguin & Tomasello 1993, Tomasello et al., 1996). In one series of experiments, Akhtar et al. (1996) created a naturalistic novel word learning situation that consisted of a sequence of events. First, the child (a 24 month old) along with three adults was introduced to three novel objects. The objects were played with successively until the child was highly familiar with each of them; however, none of these objects was ever named. Secondly, the three now familiar objects were placed in a transparent box along with a novel fourth object, the target object. The adults looked generally at the transparent container (but not at any individual object) and said, "Look, it's a modi. A modi." During the third event, the child and adults played successively with each of the four objects but the adults did not name any of them. The fourth event was the test: the four objects were placed on table and the child was asked to indicate "the modi." Although, all four objects had been present when the novel name was supplied, the children had no trouble determining the referent. The children chose the target object (the most novel of the four when the name was first supplied) when asked to get "the modi". This result fits the idea that in the moment of hearing a novel name children link that name to the object that most demands attention. The object most demanding of attention at a particular moment, in turn, depends on its novelty which depends on prior events.

A second experiment by Akhtar et al. demonstrated the role of context in the sequence of events organizing attention. First, the child and three adults played successively with three novel objects, thus making all these objects familiar to the child. Secondly, two of the adults left the room and a novel fourth object, the target object, was introduced. The child and the remaining adult played with this novel toy but the remaining adult never named the object. By one description, then, one that centres only on the object and not the context, the target object is at this point equivalent to the other three in familiarity for the child. Thirdly, the three original toys and the target object were placed together in the transparent box and the two other adults returned. While looking generally at the transparent box, the two returning adults said "Look, a gazzer. It's a gazzer." Fourthly, all participants played with the four objects. Fifthly, on the test trial, the four objects were presented and the child was asked to indicate "the gazzer". The children chose the target object.

These results fit what we know about the contextual nature of memories; with the fact that what is remembered depends on the holistic match between the general context of the original event and the context of the moment. This fact suggests the following description of events in Akhtar et al.'s second experiment. In Event 1, three memories are formed, each consist of the object (Oi) and the context of three adults playing and attending (C1): the memories formed thus are O1 + C1, O2 + C1, O3 + C1. Event 2 consists of a different object, the target OT, and a different context, one that does not contain two of the adults. Thus the memory stored of the second event will be OT + C2. In the critical Event 3 when the novel name is offered, all objects are presented in the transparent box and all four adults are present. This context, C1', is thus more similar to the original context than is C2. By this analysis, then, the target object is the most novel in the context when the name is offered and thus the one attended to most and most likely to be associated with the novel name.

There are many other results in the literature that show that the interpretation of a novel name applied to a novel object is a mental event that grows out of and is continuous with just-previous mental activity: children interpret novel words as referring to novel objects with novelty determined by just-preceding events (Markman 1989, Merriman & Bowman 1989); children are more likely to map a word to an object and form a well-organized category if the sound of the word has been highlighted by recently hearing other similar words (Merriman & Marazita 1995, Schwartz et al. 1987); and finally, presenting children with an array of test objects (and thus the variation among potential category members) prior to naming the exemplar alters the formed category (Merriman et al. 1991). Clearly, children's category formation in the context of interpreting a novel object name is a mental event created in the context of ongoing processes of perceiving and remembering.

A history of individual acts of knowing

This is considerable evidence that supports Smith's (1995) proposal that the smartness of children's category formation upon hearing a novel object derives from learned associations among linguistic contexts and the properties of novel objects. The fact that in many stimulus contexts

children form categories organized by shape similarities makes sense from this point of view. In English, the common nouns that name common concrete objects refer to categories of things of similar shape (Rosch 1978, Biederman 1987).

Jones et al. (1991) proposed that children learn to attend to shape as a consequence of learning words that refer to similar things. They reasoned that if attention to shape in the context of naming rigid objects was the product of statistical regularities in the language then it should emerge only after some number of names for rigid things have been learned. They tested this hypothesis in a longitudinal study of children from 15 to 20 months. During this time span, the number of concrete nouns in the children's productive vocabularies increased from an average of five to over 150. The children were tested once a month in an artificial word learning task much like the artificial word learning task described earlier. The principle result is that the children did not systematically generalize novel nouns to new instances by shape until after they had 50 nouns in their productive vocabulary. This result is consistent with the idea that the shape bias in naming is a product of learning words that refer to objects of similar shape.

The importance of the child's personal history of forming categories is also seen in the growing evidence for cross-linguistic differences (Imai & Gentner in press; Gopnik et al. in press; Waxman et al. in press; Choi & Bowerman 1991, Sera et al. 1991). Several of these cross-linguistic studies focused on children's naming of objects versus substances. Lucy (1992) predicted cross-linguistic differences in the categorization of these two kinds because languages differ in how they mark countable and non-countable entities – in what kinds of nouns called "count nouns" take the plural and numerical determiners (e.g. "two dogs") and what kinds called "mass nouns" do not (e.g. "some sugar"). English maps the distinction between count nouns and mass nouns on to that between rigid objects versus substances (see Lucy 1992). The work of Soja et al. (1991) cited earlier suggests that very young children learning English generalize names for objects and substances differently – a result that suggests that the object–substance distinction could be universal and independent of language.

However, other languages divide up countable and non-countable things differently and in ways that ignore the distinction between rigid things of constant shape and malleable substances of variable shape. For example, in some languages only categories referring to humans (brother, wife, priest) receive plural marking. In others, the division between kinds of entities that are syntactically marked or not marked for number is between animate things and inanimate things. Lucy specifically argued that the potency of the contrast between rigidly-shaped things and substances might be specific to languages like English which map the count–mass distinction on this contrast; languages which map the distinction on, for example, animates versus inanimates might not direct attention away from the contrast between rigid and non-rigid things.

To test this idea, Lucy (1992) examined the classifications of adult speakers of Yucatec Mayan, a language in which nouns for all inanimate things (rigid objects and non-solid substances) are treated like English mass nouns and do not take the plural. Consistent with Lucy's prediction, Yucatec-speaking adults, in marked contrast to English-speaking adults, classify rigid objects by material and not by shape. For example, given a cardboard box, a wooden box, and a piece of irregularly shaped cardboard, Yucatec speakers classify the cardboard box with the piece of cardboard; English speakers classify the two boxes together. More recent evidence (Lucy 1996) suggests further that these differences in classification between English-speaking and Yucatec-speaking individuals increase with development – as they should if they are created by the statistical regularities in a language and the individual's history of category judgements.

Imai & Gentner (1993) provided further evidence for the developmental growth of cross-linguistic differences. Their subjects were English-speaking and Japanese-speaking children. In Japanese, like Yucatec, inanimate nouns (names for rigid things and non-rigid substances) do not take the plural; all are syntactically like English mass nouns. Imai & Gentner showed that this difference matters for how children generalize novel names to new instances. Specifically, American children formed lexical categories consistent with an object–substance distinction. Japanese children's categorizations were better described by a complex-shape versus simple-shape distinction rather than rigidity *per se*. Moreover, these differences were evident at age two but increased dramatically between the ages of two and four years, and were most marked in adults. These results demonstrate that category formation in the moment – the kind of things children assume to have the same name – is moulded by the kinds of categories that children have formed in the past.

Putting it together

Young children learn names for things rapidly. They are so adept at this learning that they seem to know the category of objects a word refers to from hearing it name a single object. The evidence reviewed above strongly fits the picture in Figure 5.9. Children's word learning is smart because children create categories on-line out of their history of experience, the transient effects of preceding activity, and the details of

the moment. The wisdom of the past is fit to the idiosyncracies of the here and now. This is a powerful idea. It means that children can form novel categories – think thoughts never thought before – that are adaptive. This can happen because the categories created on-line will be a unique mix of past and present.

How this mixture can cause both stability and variability is illustrated in one final study of children's novel word generalizations. Smith et al. (1992) examined young children's interpretations of novel nouns versus novel adjectives. The objects in all conditions were made of wood and varied in shape and colour. The colours of the named objects were realized by putting glitter in paint. In the noun condition, an exemplar was labelled with a novel count noun "is a dax"; in the adjective condition, the exemplar was labelled with a novel adjective, "is a dax one". Critically, the experiment was performed twice: once when the objects were presented under ordinary lighting conditions and once when the objects were presented in a darkened chamber under a spotlight. The effect of the spotlight was to make the glitter sparkle and the colours thus attention-grabbing.

The key result is that the categories that children form depend on both syntactic context and room illumination. Under ordinary illumination the children generalized both novel nouns and novel adjectives only to objects the same shape as the exemplar. Colour was ignored. In the spotlight condition, in contrast, children generalized novel nouns and novel adjectives differently. Children again generalized the novel noun by shape, ignoring colour. But they generalized the novel adjective on the basis of the sparkling colour.

These results show that knowledge of language and local idiosyncratic forces on attention combine in a moment to invent categories. It is unlikely that these children possessed any specific rules about adjective categories and sparkling versus non-sparkling glitter. Rather, children's variable interpretation of novel adjectives is best explained in the here-and-now mix of past learning about adjectives and the context of sparkling colours just as the stability of children's interpretations of novel nouns is explained by the mix of past learning about nouns and the present task context. In brief, the processes that make stability and variability may be the same.

WHAT ABOUT CONCEPTS?

> Experience is remolding us every moment, and our mental reaction on every given thing is really a resultant of experience of the whole world up to that date. (William James 1890:234).

Children's category formation in the context of hearing a novel object named is sometimes considered apart from categorization in adults and as reflecting mechanisms perhaps specific to language learning (see Markman 1989). However, category formation by children in the service of word learning looks very much like categorization by humans generally. Wise stability is coupled with street-smart variability. The stability and the variability seem to directly fall out of basic truths about processes of perceiving and remembering: their context-dependency, the extended duration of real-time processes, and finally and most critically, that these are processes which change themselves by their own activity. Moments of knowing, be they interpreting a novel word, reasoning about quarters and pizzas, judging the essential properties of water – may be wisely stable and smartly variable all at once precisely because each moment of knowing combines intelligence across multiple timescales – the fast time of on-line processes, the slower time of developmental change, and in the non-changing aspects of perceptual and memorial processes, the very slow scale of change over evolutionary time. A successful theory of categories thus might require that we give up timeless abstractions such as concepts.

One might ask whether it is, nonetheless, useful to keep the idea of concept, just alter its meaning. Barsalou (1993) proposed such a new conception of concepts, one in which the word concept is used to refer to not to internal representations but to the thought made in the moment. James (1890:246) took the opposite stand, "I shall avoid the use of the expression concept altogether".

REFERENCES

Akhtar, N., M. Carpenter, M. Tomasello, 1996. The role of discourse novelty in early word learning. *Child Development*, **67**(2), 635–645.

Anstis, S. M. & S. Saida 1985. Adaptation to auditory streaming of frequency-modulated tones. *Journal of Experimental Psychology: Human Perception and Performance* **11**, 257–71.

Armstrong, S. L., L. R. Gleitman, H. Gleitman 1983. What some concepts might not be. *Cognition* **13**, 263–308.

Ashby, F., M. Gregory, W. Todd 1993. Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology* **37**, 372–400.

Baldwin, D. A. 1991. Infant's contribution to the achievement of joint reference. *Child Development* **62**, 857–90.

Baldwin, D. A. & E. M. Markman 1989. Establishing word-object relations: a first step. *Child Development* **60**, 174–82.

Banks, W. P. & J. Flora 1977. Semantic and perceptual processes in symbolic comparisons. *Journal of Experimental Psychology: Human Perception and Performance* **3**, 278–90.

Barsalou, L. W. 1983. Ad hoc categories. *Memory and Cognition* **11**, 211–27.

Barsalou, L. W. 1987. The instability of graded structure: implications for the nature of concepts. In *Concepts and conceptual development*, U. Neisser (ed.), 101–40. Cambridge: Cambridge University Press.

Barsalou, L. W. 1989. Intraconcept similarity and its implications for interconcept similarity. In *Similarity and analogical reasoning*, S. Vosniadou & A. Ortony (eds), 76–121. Cambridge: Cambridge University Press.

Barsalou, L. W. 1993. Challenging assumption about concepts [commentary]. *Cognitive Development* **8**, 169–80.

Bauer, P. J. & R. Fivush 1992. Constructing event representations: building on a foundation of variation and enabling relations. *Cognitive Development* **7**, 381–401.

Biederman, I. 1987. Recognition by components: a theory of human image understanding, *Psychological Review* **94**, 115–47.

Brooks, L. R. 1987. Decentralized control of categorization: the role of prior processing episodes. In *Concepts and conceptual development*, U. Neisser (ed.), 141–74. Cambridge: Cambridge University Press.

Butler, J. & C. Rovee-Collier 1989. Contextual gating of memory retrieval. *Developmental Psychobiology* **22**, 533–52.

Cech, C. G. & E. J. Shoben 1985. Context effects in symbolic magnitude comparisons. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **11**, 299–315.

Cech, C. G., E. J. Shoben, M. Love 1990. Multiple congruity effects in judgements of magnitude. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **16**, 1142–52.

Choi, S. & M. Bowerman 1991. Learning to express motion events in English and Korean: the influence of language-specific lexicalization patterns. *Cognition* **41**, 83–121.

Coren, S. & C. Porac 1977. Fifty centuries of right-handedness: the historical record. *Science* **198**, 631–2.

Eich, E. 1985. Context, memory, and integrated item/context imagery. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **11**, 764–70.

Fodor, J. D. 1977. *Semantics: theories of meaning in generative grammar*. New York: Crowell.

Gelman, S. A. & J. D. Coley 1991. Language and categorization: the acquisition of natural kind terms. In *Perspectives on language and thought: interrelations in development*, S. A. Gelman & J. P. Byrnes (eds), 146–96. Cambridge: Cambridge University Press.

Gibson, J. J. 1933. Adaptation, after effect and contrast in the perception of curved lines. *Journal of Experimental Psychology* **16**, 1–31.

Godden, D. & A. Baddeley 1980. When does context influence recognition memory? *British Journal of Psychology* **71**, 99–104.

Goldstone, R. L. 1995. Effects of categorization on color perception. *Psychological Science* **6**, 298–304.

Goldstone, R. L., D. L. Medin, D. Gentner 1991. Relational similarity and the non-independence of features in similarity judgements. *Cognitive Psychology* **23**, 222–62.

Goldstone, R., D. Medin, J. B. Halberstadt, in press. Similarity in context. *Memory and Cognition*.

Gopnik, A., S. Choi, T. Baumberger, in press. Cross-linguistic differences in early semantic and cognitive development. *Cognitive Development*.

Halff, H. M., A. Ortony, R. C. Anderson 1976. A context-sensitive representation of word meanings. *Memory and Cognition* **4**, 378–83.

Halpern, D. F. & J. S. Warm 1984. The disapperance of dichoptically presented real and subjective contours. *Bulletin of the Psychonomic Society* **22** (5) 433–6.

Harris, M. & M. Coltheart 1986. *Language processing in children and adults: an introduction*. London: Routledge & Kegan Paul.

Hasher, L. & R. T. Zacks 1984 Automatic processing of fundamental information: the case of frequency of occurrence. *American Psychologist* **39**, 1372–88.

Imai, M. & D. Gentner, 1993. Linguistic relativity vs universal ontology: cross-linguistic studies of the object substance distinction. *Proceedings of the Chicago Linguistic Society* 1993.

Imai, M., D. Gentner, N. Uchida 1994. Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive Development* **9**, 45–75.

Inhelder, B. & J. Piaget 1964. *The early growth of logic in the child*. New York: Norton.

Jacoby, L. L. 1983. Perceptual enhancement: persistent effect of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **9**, 21–38.

Jacoby, L. L., C. Kelley, J. Brown, J. Jasechko 1989. Becoming famous overnight: limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology* **56**, 326–38.

James, W. 1950. *The principles of psychology*, vol. 1. New York: Dover Publications [original work published 1890. New York: Henry Holt & Co.

Jones, S. S. & L. B. Smith 1993. The place of perceptions in children's concepts. *Cognitive Development* **8**, 113–40.

Jones, S. S., L. B. Smith, B. Landau 1991. Object properties and knowledge in early lexical learning. *Child Development* **62** (3), 499–516.

Katz, J. 1972. *Semantic theory*. New York: Harper & Row.

Keil, F. C. 1989. *Concepts, kinds, and cognitive development*. Cambridge, Mass.: MIT Press.

Keil, F. C. 1994. Explanation, association, and the acquisition of word meaning. *Lingua* **92**, 169–96.

Kelly, M. H. & S. Martin 1994. Domain-general abilities applied to domain-specific tasks: Sensitivity to probabilities in perception, cognition, and language. *Lingua* **92**, 105–40.

Klatzky, R. L. 1980. *Human memory: structures and processes*, 2nd edn. San Francisco: W. H. Freeman.

Komatsu, L. D. 1992. Recent views of conceptual structure. *Psychological Bulletin* **112**, 500–26.

Lakoff, G. 1987 Cognitive models and prototype theory. In *Concepts and conceptual development*, U. Neisser (ed.), 63–100. New York: Cambridge University Press.

Landau, B., L. B. Smith, S. S. Jones 1988. The importance of shape in early lexical learning. *Cognitive Development* **3**, 299–321.

Landau, B., L. B. Smith, S. Jones 1995. *Object shape, object function and object name*.

Lewicki, P., T. Hill, I. Sasaki 1989. Self-perpetuating development of encoding biases. *Journal of Experimental Psychology: General* **118**, 323–37.

Lichtenstein, S., P. Slovic, B. Fischoff, M. Layman, B. Combs 1978. Judged frequency of lethal events. *Journal of experimental Psychology: Human Learning and Memory* **4**, 551–78.

Light, L. L. & L. Carter-Sobel 1970. Effects of changed semantic context on recognition memory. *Journal of Verbal Learning and Verbal Behavior* **9**, 1–11.

Lucy, J. A. 1992. *Grammatical categories and cognition: a case study of the linguistic relativity hypothesis*. Cambridge: Cambridge University Press.

Lucy, J. A. 1996. The impact of language-specific categories on the development of classification behaviour. *International Journal of Psychology* **31** (3–4), 5623.

Malt, B. C. 1994. Water is Not H_2O. *Cognitive Psychology* **27**, 41–70.

Markman, E. M. 1989. *Categorization and naming in children: problems of induction*. Cambridge, Mass.: MIT Press.

Marks, L., E. 1993. Contextual processing of multidimensional and undimensional auditory stimuli. *Journal of Experimental Psychology Human Perception and Performance* **19**, 227–49.

McGraw, G. & J. Rehling 1994. *Roles in letter perception: human data and computer models*.

Medin, D. & A. Ortony. 1989. Psychological essentialism. In *Similarity and analogical reasoning*, S. Vosniadou & A. Ortony (eds), 179–95. New York: Cambridge University Press.

Merriman, W. E. & L. L. Bowman 1989. *The mutual exclusivity bias in children's word learning*. Monograph 54 (p. 130), Society for Research in Child Development.

Merriman, W. E. & J. M. Marazita 1995. The effect of hearing similar-sounding words on young 2-year-olds' disambiguation of novel noun reference. *Developmental Psychology* **31**, 973–84.

Merriman, W. E., J. M. Schuster L. B. Hager 1991. Are names ever mapped onto preexisting categories? *Journal of Experimental Psychology: General* **120**, 288–300.

Mervis, C. B. 1987. Child-basic object categories and early lexical development. In *Concepts and conceptual development*, U. Neisser (ed.), 201–233. Cambridge: Cambridge University Press.

Minsky, M. 1986. *The society of mind*. New York: Simon & Schuster.

Murphy, G. L. & D. L. Medin 1985. The role of theories in conceptual coherence. *Psychological Review* **92**, 289–316.

O'Leary, A. & M. McMahon 1991. Adaptation to form distortion of a familiar shape. *Perception and Psychophysics* **49**, 328–32.

Olguin, R. & M. Tomasello 1993. Twenty-five-month-old children do not have a grammatical category of verb. *Cognitive Development* **8**, 245–72.

Palmer, S. G. 1989. Reference frames in the perception of shape and orientation. In *Object perception: structure and process*, B. E. Shepp & S. Ballesteros (eds), 121–64. Hillsdale, NJ: Erlbaum.

Palmeri, T. J., S. D. Goldinger, D. B. Pisoni 1993. Episodic encoding of voice attributes and recognition memory for spoken words. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **19**, 309–28.

Perris, E. E., N. A. Myers R. K. Clifton 1990. Long-term memory for a single infancy experience. *Child Development* **61**, 1796–1807.

Reber, A. S., S. M. Kassin, S. Lewis G. Cantor 1980. On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning and Memory* **6**, 492–502.

Regier, T. 1995. A model of the human capacity for categorizing spatial relations. *Cognitive Linguistics* **6**, 63–88.

Remez, R. E. 1979. Adaptation of the category boundary between speech and nonspeech: a case against feature detectors. *Cognitive Psychology* **11**, 38–57.

Remez, R. E., J. E. Cutting, M. Studdert-Kennedy 1980. Cross-series adaptation using song and string. *Perception and Psychophysics* **27**, 524–30.

Rips, L. J. 1989. Similarity, typicality, and categorization. In *Similarity and analogical reasoning*, S. Vosniadou & A. Ortony (eds), 21–59. Cambridge: Cambridge University Press.

Rips, L. J. & A. Collins 1993. Categories and resemblance. *Journal of Experimental Psychology: General* **122**, 468–86.

Rosch, E. 1973. On the internal structure of perceptual and semantic categories. In *Cognitive development and the acquisition of language*, T. E. Moore (ed.), 111–44. San Diego, California: Academic Press.

Rosch, E. 1978. Principles of categorization. In *Cognition and categorization*, E. Rosch & B. Lloyd (eds), 28–46. Hillsdale, New Jersey: Erlbaum.

Rosch, E. & C. B. Mervis 1975. Family resemblances: studies in the internal structure of categories. *Cognitive Psychology* **7**, 573–605.

Rosenbaum, D. A. 1991. *Human motor control*. San Diego, CA: Academic Press.

Rovee-Collier, C., P. C. Griesler, L. A. Earley 1985. Contextual determinants of retrieval in three-month-olds infants. *Learning and Motivation* **16**, 139–57.

Saegert, S., W. Swap, R. B. Zajonc 1973. Exposure, context, and interpersonal attractors. *Journal of Personality and Social Psychology* **25**, 234–42.

Salasoo, A., R. M. Shiffrin, T. C. Feustel 1985. Building permanent memory codes: codification and repetition effects in word identification. *Journal of Experimental Psychology: General* **114**, 50–77.

Sanocki, T. 1991. Intra- and inter-pattern relations in letter recognition. *Journal of Experimental Psychology: Human Perception and Performance* **17**, 924–41.

Sanocki, T. 1992. Effects of font- and letter-specific experience on the perceptual processing of letters. *American Journal of Psychology* **105**, 435–58.

Schwartz, R. G., L. B. Leonard, D. M. Frome-Loeb, L. A. Swanson 1987, Attempted sounds are sometimes not: an expanded view of phonological selection and avoidance. *Journal of Child Language* **14**, 411–18.

Schyns, P. G. & L. Rodet, in press. Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory & Cognition.*

Sera, M. D., C. A. H. Berge, J. D. C. Pintado 1994. Grammatical and conceptual forces in the attribution of gender by English and Spanish speakers. *Cognitive Development* **9**, 261–92.

Sera, M. D., E. L. Reittinger, J. D. C. Pintado 1991. Developing definition of objects and events in English and Spanish speakers. *Cognitive Development* **6**, 119–42.

Sera, M. D., D. Troyer, L. B. Smith 1988. What do two-year-olds know about the sizes of things? *Child Development* **59**, 1489–96.

Shapiro, B. J. 1969. The subjective estimation of relative word frequency. *Journal of Verbal Learning and Verbal Behaviour* **13**, 638–43.

Shedler, J. K., J. Jonides 1985. *Availability: plausible but questionable*. Paper presented at the 26th Annual Meeting of the Psychonomic Society, Boston.

Smith, E. E. & D. L. Medin 1981. *Categories and concepts*. Cambridge, Mass.: Harvard University Press.

Smith, E. E. & S. A. Sloman 1994. Similarity- versus rule-based categorization. *Memory and Cognition* **22**, 337–86.

Smith, L. B. 1995. *Self-organizing processes in learning to learn words: development is not induction*. The Minnesota Symposia on Child Psychology, vol. 28: Basic and applied perspectives on learning, cognition, and development (1–32). Hillsdale, New Jersey: Erlbaum.

Smith L. B., S. S. Jones, B. Landau 1992. Count nouns, adjectives, and perceptual properties in children's novel word interpretations. *Developmental Psychology* **28**, 273–86.

Smith, L. B., S. Jones, B. Landau, 1996. Naming in young children: a dumb attentional mechanism. *Cognition* **60**, 143–171.

Smith, S. M. 1986. Environmental context-dependent recognition memory using a short-term memory task for input. *Memory and Cognition* **14**, 347–54.

Soja, N. 1992. Inferences about the meanings of nouns: the relationship between perception and syntax. *Cognitive Development* **7**, 29–46.

Soja, N., S. Carey, E. Spelke 1991. Ontological categories guide young children's inductions of word meanings: object terms and substance terms. *Cognition* **38**, 179–211.

Tomasello, M. & A. C. Kruger 1992. Joint attention on actions: acquiring verbs in ostensive and non-ostensive contexts. *Journal of Child Language* **19**, 311–23.

Tomasello, M., R. Strosberg, N. Akhtar, 1996. Eighteen-month-old children learn words in non-ostensive contexts. *Journal of Child Language* **23**, 157–176.

Tulving, E. & D. M. Thomson 1973. Encoding specificity and retrieval processes in episodic memory. *Psychological Review* **80**, 352–73.

Waxman, S. R., A. Senghas, S. Benveniste, in press. *A cross-linguistic examination of the noun-category bias: evidence from French- and Spanish-speaking preschool-aged children*.

Wittgenstein, L. 1953. *Philosophical investigations*. New York: Macmillan.

Zatorre, R. B. 1968. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, Monograph Supplement* **9**, 1–29.

Zatorre, R. J. & A. R. Halpern 1979. Identification, discrimination, and selective adaptation of simultaneous musical intervals. *Perception and Psychophysics* **26**, 384–95.