CHAPTER 13

Concepts and Categorization

Douglas L. Medin
John D. Coley

I. INTRODUCTION

It has been said that psychology has a long past but a short history. Nowhere is this observation more apt than in the psychology of concepts. The study of concepts and categorization is directed at the most fundamental questions concerning the interaction of mind and world. What kinds of things are there in the world and how do people come to know them? What is the relationship between lexical terms and their referents? Do we simply *recognize* categories or do constructive processes operate to *create* categories and concepts? These are questions asked by early Greek philosophers, questions that persisted in monasteries through the dark ages, and questions that organized thought in biology from the emergence of systematics in the 17th century, through the development of evolutionary theory, and they continue to the present day.

Psychology and experimental studies of concepts and categorization, while burgeoning, are still in their infancy. As such, one might expect false steps, some falls, and uneven progress. As we shall see, the psychological study of concepts has lived up to this metaphor (to the occasional frustration of researchers). But learning to walk or even crawl promises new views and perspectives. In this chapter, we provide a selective review of research on concepts and categorization, focusing on the last 30 years or so. Even in this short span there have been a number of twists and turns, and it is not always easy to identify genuine progress. One should note, however, that progress also comes in the form of methodological and technical advances

Perception and Cognition at Century's End

that allow researches to ask sharper questions. From the latter perspective, the field has seen unambiguous advances.

An important characteristic of categorization research is a multiplicity of approaches. Although pluralism can be taken as an index of a lack of clear goals and means, we prefer to read it as a "divide and conquer" strategy appropriate for problems of the magnitude of the psychology of concepts. At the same time, however, there needs to be enough communication to recognize convergences and exploit progress in allied subareas. In that spirit our review will indicate points where greater interplay across approaches may be fruitful.

The remainder of this chapter is organized as follows. We first continue our somewhat belabored introduction by outlining the sometimes fitful progress in the psychological study of concepts and categories over the past 50 years or so. Next we turn to constraints and challenges that qualify key results, generate new findings, and (in our opinion) ultimately constrain what any theory of categorization must account for. This sets this stage for a discussion of current challenges and further opportunities for progress.

II. CONCEPTUAL FUNCTIONS

First, it will be useful to have a few definitions. By *concept* we mean a mental representation of a category serving multiple functions, one of which is to allow for the determination of whether or not something belongs to the class. A *category* refers to the set of entities picked out by the concept. As we shall see, concepts serve multiple functions and a focus on a single function comes at the risk of developing theories that are too narrow to do the work they will ultimately be asked to do (e.g., Matheus, Rendell, Medin, & Goldstone, 1989).

There are numerous possible taxonomies, but we will distinguish seven functions of concepts: categorization, understanding, inferences, explanation and reasoning, learning, communication, and combination. These are meant to be neither mutually exclusive nor exhaustive, but they illustrate the vast array of cognitive functions performed by concepts, and the enormous task facing any theory of concepts and categorization.

The categorization function of concepts refers to the fact that mental representations are used to determine the category membership of entities. Categorization allows us to bring relevant knowledge to bear in the service of *understanding* and *making predictions.* For example, after categorizing some novel object as a member of the category step ladders, people can understand its relevant parts and know how to interact with it. Categorization also allows one to make predictions or *inferences* concerning the affordances of some entity. For example, one might infer that this novel step ladder would support one's weight. Concepts are critically involved in *explanation* and reasoning. Having categorized some patient as an *anorexic* (a form of eating disorder) one might be able to explain why they insist on adding skim milk rather than half and half to their coffee.

stions. From the latter perspective, the field

tegorization research is a multiplicity of taken as an index of a lack of clear goals and and conquer" strategy appropriate for probgy of concepts. At the same time, however, ion to recognize convergences and exploit our review will indicate points where greater tful.

rganized as follows. We first continue our utlining the sometimes fitful progress in the tegories over the past 50 years or so. Next hat qualify key results, generate new findonstrain what any theory of categorization a discussion of current challenges and fur-

initions. By *concept* we mean a mental reple functions, one of which is to allow for nething belongs to the class. A *category* refers oncept. As we shall see, concepts serve mulnction comes at the risk of developing thek they will ultimately be asked to do (e.g., e, 1989).

nies, but we will distinguish seven functions ing, inferences, explanation and reasoning, tion. These are meant to be neither mutulustrate the vast array of cognitive functions ous task facing any theory of concepts and

epts refers to the fact that mental represenry membership of entities. Categorization to bear in the service of *understanding* and tegorizing some novel object as a member understand its relevant parts and know how llows one to make predictions or *inferences* tity. For example, one might infer that this weight. Concepts are critically involved in gorized some patient as an *anorexic* (a form to explain why they insist on adding skim offee.

Concepts also support *learning* in that new entities are not only understood in terms of old but also feed back to modify or update contexts. For example, learning that a penguin is a bird not only adds another instance to *bird* but also may cause the learner to rethink what is meant by *bird* to begin with. Just how this updating is done is an important theoretical question. For example, category modification needs to balance the need to be sensitive to particular contexts with the danger of discarding accumulated wisdom derived from a broader range of contexts.

Obviously we also use concepts for *communication*. The interpersonal aspect of concepts places constraints on virtually every other function. Communication is facilitated to the extent that conceptions of categories are shared across language users. Furthermore, communicative goals may determine whether some entity is referred to as a *cottontail*, a *rabbit*, or simply an *animal* (Grice, 1975).

Finally, *combining concepts* allows us to express and create an unlimited number of new concepts. People are able to understand novel combinations of concepts such as *green mouse* (presumably a mouse that is green) and *green uncle* (an uncle with lots of money, or perhaps a new and inexperienced uncle).

These functions place competing demands on conceptual structure. For example, one can *categorize* most easily by focusing attention solely on the presence or absence of a single salient feature (e.g., whether or not something is red). But the category *red things* fails to support other conceptual functions; knowing that something is a member of *red things* does little to support inferences about the object, increase understanding, or provide explanations. Thus, different conceptual functions might make different and even competing demands on conceptual structure.

III. THE EMPIRICAL STUDY OF CATEGORIZATION

We do not need to rehearse here the general merits of controlled experimental contrasts. Instead, our aim is to call attention to the ongoing interaction between laboratory studies of conceptual behavior and so-called real-world categories. The general form of this interaction has been as follows. On the basis of observation, intuition, or argument, ideas about the structural underpinnings of natural (real-world) categories are developed. Associated with these conjectures may be some ideas or theories concerning both the mental representations people develop of such categories and the processes that create and operate on these representations. To test these ideas, researchers often construct "artificial categories" in the laboratory where the "structure" is unambiguous or independently established. Participants in experiments then learn these categories and then may be given various tests of their knowledge. The resulting performances are used to test ideas and theories concerning representation and processing. Often these studies with artificial categories are run in parallel with studies of natural categories. If the experimental findings are also parallel, researchers have increased confidence that they have, in fact, successfully brought relevant structural properties into the laboratory. To the extent that this goal is realized, one also has more confidence that the laboratory findings with respect to representation and processing will generalize to real-world (natural) categories. Lack of corresponding findings suggests either that the relevant structure has not been captured or that conditions of category learning and use in the world differ from those in the laboratory. Under these circumstances, findings with natural and with artificial categories each may be useful in their own right, but questions about generalizing from artificial to natural categories loom large.

In short, analyses of natural categories and experiments with artificial categories are closely linked. To be sure, researchers will sometimes explicitly depart from realistic category structures in order to set up some contrast to theoretical interest. But even here the researcher does not lose sight of real-world categories. Explicit departure requires an awareness of "from what," and does not constitute a neglect of structure.

Why are we hammering away on this point about relationships between the natural and the artificial? One reason is that it's easy to get lost in the details of laboratory contrasts and to interpret differences in the trade-off between artificiality and generalizability as representing incompatible approaches to the study of categorization. We prefer to see the tension between laboratory control and worldly realism as admitting a variety of research strategies that may be variously successful depending on one's goals. But what is common to them all is that no one believes that artificial categories are an end in themselves, and everyone would like to make correct generalizations to and inferences about natural categories.

IV. SIMILARITY

Everyone agrees that category membership is not arbitrary, not simply a list where the only thing category members have in common is that they are on the list. For example, we put salmon in a category with trout rather with robins or pencils (and if the choices were restricted to the latter two, we would pick robins over pencils). The general claim is that the basis for this categorization goes beyond the simple knowledge that salmon and trout are *fish*. But precisely what *is* relevant to placing salmon and trout in the same category?

One very important idea is that *similarity* is the organizing principle for categories and categorization. That is, *salmon* fits into a category with *trout* because salmon are more similar to trout than they are to robins or pencils. Entities in the same category are generally more similar to each other than they are to examples from contrasting categories. Certainly this principle makes intuitive sense and it has been rarely contested (we'll consider an exception a bit later). Instead, the debate has concerned how to analyze the notion of similarity and, in turn, how to map this analysis onto principles of category membership.

The modal approach to analyzing similarity has been to assume that concepts are comprised of features. Two entities are similar to the extent that they share (underlying) features (e.g., Tversky, 1977). That is, the molar notion of similarity is to be understood in terms of molecular processes of feature matches and mismatches. For example, one's mental representation of *salmon* and *trout* may include the features of having scales, gills, and fins and being alive. For *robin* the corresponding features might be having feathers, lungs, wings, and being alive. We'll defer the question of what features of pencils correspond to gills or lungs and simply note that pencils are not alive. By this featural approach, salmon are very similar to trout because they share four features and more similar to robins than to pencils because of the shared feature of being alive.

One nice property of featural analyses is that they integrate perceptual and conceptual similarity. There is evidence that the nervous system has "feature detectors" sensitive to perceptual properties, and the idea is to generalize this approach to the conceptual level. On this view, even abstract concepts like *liberty* can be decomposed into features and compared with other concepts like *freedom* and *justice.* As we'll see, a critical question is just how to perform a decomposition of concepts into features. Strategies have ranged from formal semantic analysis to simply asking people to list things that are true of concepts.

With this abstract description as background we are ready to turn to the question of how this analysis of similarity into features has been linked to the empirical study of category structure. The next section will briefly describe the two main views of conceptual structure that have organized research on concepts (see Smith & Medin, 1981, for an extended review).

V. CATEGORY STRUCTURE

Much of the research on concept learning in the third quarter of this century had as its primary focus learning (e.g., Bourne, 1970; Levine, 1975; Trabasso & Bower, 1968). An important issue concerned whether learning was all-or-none or gradual, and the upshot of these studies was support for all-or-none learning. The corresponding theoretical analyses used hypothesis testing as a framework. The idea was that category learning consisted of testing hypotheses or rules for category membership—incorrect hypotheses would lead to chance performance, but as soon as the learner tried out the correct rule or hypothesis, categorization performance would be perfect.

A. The Classical View

Although, to our knowledge, no one was at that time drawing explicit connections between these rule-learning experiments and the structure of natural categories, the two orientations were compatible. Specifically, the prevailing idea was that natural categories were structured in terms of singly necessary and jointly sufficient features (Katz & Postal, 1964). This has come to be called the classical view of concepts. A set of necessary and sufficient features means that category membership is determined by a conjunctive rule. If some entity has the set of requisite features, its a member of the category; otherwise it is not. Consider, for example, the concept *triangle.* A triangle is a closed geometric form with three sides and interior angles that sum to 180 degrees. If any one of these properties is missing, one does not have a triangle.

It is easy to go from ideas about category structure to conjectures about learning. The classical view is consistent with the idea that learning is based on hypothesis testing where hypotheses concern which features are defining. When the set of defining features has been mastered one has a rule for determining category membership.

Note that the classical view falls within the general framework of similarity models. Every member of a category shares features (namely, the defining features) with every other member of the category, and nonmembers differ from members in at least one of these features. From this perspective one could predict that the more defining features a nonmember shares with category members the harder it should be to reject as a member. For example, people should take longer to say that a robin is not a trout than to say that a pencil is not a trout (see Smith, Shoben, & Rips, 1974).

The classical view seems to work for triangles, but is it correct for natural categories more generally? The consensus among psychologists who study concepts is that the classical view is inadequate as a theory of conceptual structure. This judgment is driven both by doubts about whether concepts necessarily have defining feature and by evidence that people's conceptual behavior is not restricted to defining features. The arguments and counter-arguments can become quite complex (see Smith & Medin, 1981, for details) but the most serious problems with the classical view can be readily summarized.

First, if concepts have defining features, we should be able to say what they are. Yet even seemingly simple concepts like *dog* or *game* defy analysis into defining features. People may believe that concepts have defining features, but the features given as candidates may not hold up to closer scrutiny (McNamara & Sternberg, 1983). For example, people may list "flies" as a necessary property of *bird* but clearly not all birds fly (ostriches, penguins, baby birds).

Of course, one might argue that defining features are not necessarily accessible to consciousness. Certain syntactic rules are inaccessible to laypeople who nonetheless follow them. Perhaps all the classical view requires is a procedure for determining category membership. But this raises a second problem—there are numerous cases in which it is not clear whether or not an example belongs to a category. Is a radio an instance of furniture? What about a rug? People not only disagree with each other concerning category membership but also show internal inconsistency when questions are asked on separate occasions (Barsalou, 1989; Bellezza, 1984; McCloskey & Glucksberg, 1978).

One might argue that uncertainty about category membership just reflects uncertainty about whether some necessary feature is present. Even when category membership is certain, however, some examples of a concept seem to be better than others. For example, people judge a robin to be a better example of *bird* than a turkey 'y

: form with three sides and interior angles .ese properties is missing, one does not have

:gory structure to conjectures about learn- ı the idea that learning is based on hypoth- /hich features are defining. When the set of has a rule for determining category mem-

n the general framework of similarity mod- 'eatures (namely, the defining features) with ıd nonmembers differ from members in at :rspective one could predict that the more /ith category members the harder it should ›eople should take longer to say that a robin is not a trout (see Smith, Shoben, & Rips,

: triangles, but is it correct for natural cate- ımong psychologists who study concepts is theory of conceptual structure. This judg- vhether concepts necessarily have defining nceptual behavior is not restricted to defin- -arguments can become quite complex (see he most serious problems with the classical

res, we should be able to say what they are. : *dog* or *game* defy analysis into defining fea- ıave defining features, but the features given : scrutiny (McNamara & Sternberg, 1983). a necessary property of *bird* but clearly not irds).

:ining features are not necessarily accessible are inaccessible to laypeople who nonethe- l view requires is a procedure for determin- es a second problem—there are numerous r not an example belongs to a category. Is a bout a rug? People not only disagree with ership but also show internal inconsistency occasions (Barsalou, 1989; Bellezza, 1984;

about category membership just reflects sary feature is present. Even when category xampies of a concept seem to be better than n to be a better example of *bird* than a turkey

is and in a speeded categorization task are faster at verifying the category membership of good examples than for poor examples (e.g., Smith et al., 1974; Rosch & Mervis, 1975).

Again one might rescue the classical view by suggesting that some features help to determine that other (defining) features are present. Good examples may have more clues to the presence of defining features, and this may account for goodness of example judgments and category verification times. Although the above ways of salvaging the classical view are not wildly implausible, they have the effect of insulating it from psychological data. In effect, the classical view becomes less relevant and less useful in organizing work on categorization. Even this might not be fatal were it not for the fact that an alternative view of conceptual structure became prominent, a view capable of addressing the categorization phenomena we have just been discussing.

B. The Probabilistic View

The probabilistic view, true to its name, claims that concepts are organized in terms of properties or features that are only characteristic of category instances. For example, *flies* may be a feature of the concept *bird* because most (but not all) birds fly. An example belongs to a category if it has enough of these characteristic features. Whatever the characteristic features are of the category *bird,* birds have more of these features than nonbirds. Certain features may be necessary, and therefore weighted heavily, but probabilistic features (usually but not always present in category members) also influence categorization (see also Cutting, chap. 4, this volume; Proffitt & Kaiser, chap. 7, this volume).

The probabilistic view was motivated by a series of important studies by Eleanor Rosch. Consider, for example, the work reported in Rosch and Mervis (1975). In one condition participants were asked to list features of category examples. In another participants listed superordinates of category members and then listed features of members of contrasting categories. The measures were then correlated with goodness of example or typicality rating. Typical (or good) members of a category tended to share features with other category members and not to share features of members of contrasting categories. Atypical or poor category members showed the opposite patterns: They were less likely to share features with category members and more likely to share features with members of contrasting categories. In short, the feature-listing data predicted goodness-of-example ratings with considerable precision.

Of course, ratings and feature listings are both dependent variables, and so far the data are correlational. But Rosch and Mervis went beyond this by creating artificial categories patterned after the feature-listing data. Specifically, they created analogs of natural categories where the examples were strings of letters and numbers. Each letter or number was considered to be a feature, and each example had five such features. The distribution of features within and across categories paralleled the feature-listing data from the early experiments. Individual examples had high, medium, or low family resemblance scores based on feature distribution. That is, examples with high family resemblance scores shared feature with other members of their own category and tended not to share features with members of contrasting categories. Participants learned the categories to a criterion, then received speeded category verification tests and finally made typicality ratings. Rosch and Mervis found that examples with high family resemblance scores were easiest to learn, had the fastest category verification times, and were rated as highest in typicality. Overall then, the studies with both artificial and natural categories converged to support the probabilistic view.

The probabilistic view is consistent both with goodness-of-example effects and unclear cases. Robins may be better examples of the category *bird* than turkeys because they have more characteristic features of birds than turkeys (e.g., singing, eating worms, migrating). Unclear cases may arise because "having enough characteristic features" is a criterion that examples may come very close to. The criterion itself might vary somewhat across occasions and lead to individual variability in category membership decisions.

This view is a prototypical similarity-based approach to categorization. If an example has more characteristic features of category A than category B (more similar to A than B) then it is placed in category A. Of course, one is still left with issues such as how one determines characteristic features and a criterion for classification but investigators have employed a variety of converging techniques that support the probabilistic view (e.g., see Mervis & Rosch, 1981, for a review).

One consequence of the probabilistic view is the idea that category membership is determined by computing similarity to a prototype. If category membership is probabilistic, then rule-based classification procedures should be less effective. If category membership is based on features that are only generally true of a category then a hypothesis-testing strategy that abandons a rule in the face of negative evidence will not be likely to succeed in learning. As an alternative instantiation of the probabilistic view, people have suggested that, based on experience with exemplars, people form an impression of the central tendency (e.g., mean or modal values for features) and that category judgments come to be based on this central tendency or prototype.

Other work with artificial categories conducted around this time provided evidence consistent with prototype formation. In these studies one begins with some prototype pattern, such as a meaningless pattern of dots, and generates category examples to be used in training by distorting the prototype pattern to varying degrees (e.g., by shifting the locations of the dots; see Posner & Keele, 1968, 1970). After training old patterns, new distortions and the prototype are presented to be classified. One striking result is that participants are as good or better at classifying the prototype pattern that they have never seen before, as they are at categorizing old patterns (e.g., Homa & Chambliss, 1975; Peterson, Meagher, Chait, & Gillie, 1973). Transfer performance is well predicted by distance from the prototype, show-

:xperiments. Individual examples had high, ores based on feature distribution. That is, : scores shared feature with other members 'o share features with members of contrast- e categories to a criterion, then received finally made typicality ratings. Rosch and . family resemblance scores were easiest to on times, and were rated as highest in typi- h artificial and natural categories converged

ɔoth with goodness-of-example effects and xamples of the category *bird* than turkeys 'eatures of birds than turkeys (e.g., singing, :s may arise because "having enough char- amples may come very close to. The crite- occasions and lead to individual variability

:y-based approach to categorization. If an ; of category A than category B (more sim- ;ory A. Of course, one is still left with issues tic features and a criterion for classification y of converging techniques that support the .osch, 1981, for a review).

ic view is the idea that category member- :ity to a prototype. If category membership ation procedures should be less effective. If es that are only generally true of a category bandons a rule in the face of negative evi- rning. As an alternative instantiation of the l that, based on experience with exemplars, il tendency (e.g., mean or modal values for ›me to be based on this central tendency or

conducted around this time provided evi- ion. In these studies one begins with some ss pattern of dots, and generates category storting the prototype pattern to varying ` the dots; see Posner & Keele, 1968, 1970). ons and the prototype are presented to be :icipants are as good or better at classifying ver seen before, as they are at categorizing 1975; Peterson, Meagher, Chait, & Gillie, .cted by distance from the prototype, show-

ing a typicality or goodness-of-example effect. Finally, when delays on the order of days are inserted between learning and transfer tests, significantly greater forgetting is observed for the old training stimuli than for the prototype or other new patterns (e.g., Homa, Cross, Cornell, Goldman, & Schwartz, 1973; Posner & Keele, 1970). This result is consistent with the idea that judgments are based on a mixture of specific items and category level (prototype) informance and that the specific item information is forgotten more rapidly. As the delay increases, by this account, it is increasingly likely that judgments will be based on the prototype. All together this seems to add up to strong parallels between natural and artificial categories and strong evidence in favor of prototype theory.

There has not been a great deal of discussion of the exact form of prototype representations (but see Farah & Kosslyn, 1982; Barsalou, 1993), but even at this abstract level of description, it has been straightforward to contrast prototype theory with alternative theories. But we are getting ahead of ourselves.

Another important consequence of the probabilistic view was a seminal series of studies by Rosch, Mervis, and their associates (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) looking at the hierarchical structure of natural categories. Objects and events in the world can be categorized at a number of different levels, varying in their abstractness. For example, a *large mouth bass* is also a *bass,* a *fish,* a *vertebrate,* an *animal,* and a *living thing.* Rosch et al. singled out one such level in such hierarchies, which they called the *basic level,* as playing a central role in categorization. Let's take a closer look at what they did. In a feature-listing task participants are able to list many features in common for members of categories like *chair, car,* and *dog* and list many fewer common features for superordinate categories like *furniture, vehicle,* and *animal.* In addition, there is only a small increase in common features listed for subordinate categories, such as *recliner, convertible,* and *poodle.* Rosch et al. (1976) took these feature-listing data as suggesting that the level of *chair, car,* and *dog* was the most informative level of categorization; that is, the basic level. Furthermore, Rosch et al. used a variety of converging operations and they each converged on the same level as basic. For example, the basic level is preferred in naming, first learned by children, the most abstract level at which an averaged (prototype) shape can be recognized, and the level at which people can categorize most rapidly.

These observations (which we can only call stunning) were interpreted as indicating that entities in the world come in natural chunks or clusters of correlated features (e.g., animals with feathers are also likely to have wings and beaks), and that human cognition is sensitive to these chunks. The clusters maximize within-category similarity relative to between-category similarity at the basic level according to Rosch et al. (1976). They did suggest that experts might become sensitive to additional features and that subordinate categories might become basic, but the central message was that (basic level) categories are given by the structure of things in the world.

Again analogs to studies of the basic level have been conducted with artificial categories (see Lassaline, Wisniewski, & Medin, 1992, for a review), but here we must sound our first discordant note. Studies with artificial categories have almost exclusively used a single measure of basicness—the level in a hierarchy where categorization is fastest. More relevant for our purposes is that the artificial categories have been structured to have a defining feature or features at the basic level, a practice in line with the Classical rather than the probabilistic view. Consequently, the parallelism and generalizability of studies with artificial categories is in question (we hasten to add that these studies may be useful in developing and testing categorization models; however, their probative value for understanding basic-level categories is in doubt).

This has been a long brief history so let's pause a bit. So far our capsule history has taken us to the mid-1970s. There is little doubt that the most exciting and important development in the psychology of concepts in the 1970s was the shift from the classical to the probabilistic view, and the corresponding emergence of ideas about the basic level and prototype representations. A critical event in this shift was the strategy of directly analyzing the structure of natural categories and bringing conjectures about structure into controlled laboratory studies. We now turn to a more detailed review and analysis of research conducted either within the framework of the probabilistic view or in reaction to it.

VI. CONSTRAINTS AND CHALLENGES FOR THEORIES OF CATEGORIZATION

In this section we consider theoretical and empirical challenges to both prototype and exemplar versions of the probabilistic view, and more general constraints to any similarity-based approach to categorization.

A. Challenges for Prototype Approaches

Earlier, we argued that different conceptual functions require somewhat different kinds of information; it should not be surprising if conceptual representations contained more information than would be expected on the basis of any one function considered by itself. Indeed, virtually all theories of concept learning assume that some types of information are preserved and other types are lost or inaccessible; comparing these assumptions to data has motivated an important component of research on classification learning. As we shall see, human conceptual behavior displays conservatism with respect to category information (Medin & Ross, 1989), a conservatism that poses problems for prototype theory.

1. Discarding Information

Prototype theory implies that the only information abstracted from categories as their central tendency. A prototype representation discards information on category size, variability of examples, and correlations among attributes (e.g., large spoons are more likely to be made of wood than small spoons). Studies with both artificial and natural categories suggest, however, that people are sensitive to all three of these types of information (Billman & Knutson, 1996; Estes, 1986a; Flannagan, Fried, & Holyoak, 1986; Fried & Holyoak, 1984; Malt & Smith, 1983; Medin, Altom, Edelson, & Freko, 1982; Medin & Schaffer, 1978; Medin & Shoben, 1988). Even notions about typicality can be shown to be sensitive to context. When one reads about a bird on a Thanksgiving platter, one does not instantiate the concept *bird* with a robin or a sparrow (Roth & Shoben, 1983). In short, prototype representations appear to discard too much information that can be shown to be relevant to human categorization.

2. Constraints on Learning

Another problem for prototypes is that they make the wrong predictions about which category structures should be easy or difficult to learn. Classifying examples on the basis of their similarity to prototypes is equivalent to a summing of evidence against a criterion. Therefore, in order to succeed, prototype models require some weighted combination of properties that accepts all category members and rejects all noncategory members. The technical term for this constraint is that categories must be linearly separable (Sebestyn, 1962) if categorization by prototypes is to be successful.

Thus, prototypes require that categories are linearly separable. As such, a key question is whether linear separability facilitates human category learning. That is, with other factors controlled, people should find it easier to learn categories that are linearly separable than those that are not. If so, this would provide important support for prototype theory. This support has not, however, appeared. Studies employing a variety of instructions, stimulus materials, and subject populations have failed to find evidence that linear separability facilities human category learning (e.g., Kemler-Nelson, 1984; Medin & Schwanenflugel, 1981; see also Shepard, Hovland, & Jenkins, 1961, and Nosofsky, Gluck, Palmeri, McKinley, & Glauthier, 1994). Later on, we will take up an exception to this rule, but the failure of linear separability represents a serious problem for prototype theory.

3. Prototypes versus Exemplars

One response to the above limitations of prototype theory is to assume that people both form prototypes and store information about specific examples. The idea would be that prototypes are needed to account for the findings considered earlier where prototypes are classified more accurately than old exemplars and show less forgetting over a retention interval. Storage of old examples would be invoked to explain the problems for prototypes just mentioned. Attractive as this strategy might appear, it has not generated much support (Busemeyer, Dewey, & Medin, 1984; see also Homa, 1984; Homa, Sterling, & Trepel, 1981), and it is instructive to see why.

First of all, to posit exemplar storage is only part of the problem; one also needs

& Medin, 1992, for a review), but here we
Studies with artificial categories have almost
asicness—the level in a hierarchy where cat-
our purposes is that the artificial categories
g feature or features at the basic level, a prac-
an the probabilistic view. Consequently, the
es with artificial categories is in question (we
useful in developing and testing categoriza-
alue for understanding basic-level categories

o let's pause a bit. So far our capsule history
is little doubt that the most exciting and
logy of concepts in the 1970s was the shift
view, and the corresponding emergence of
e representations. A critical event in this shift
he structure of natural categories and bring-
ontrolled laboratory studies. We now turn to
research conducted either within the frame-
action to it.

ENGES FOR THEORIES

and empirical challenges to both prototype
tic view, and more general constraints to any
ition.

oaches

ptual functions require somewhat different
urprising if conceptual representations con-
e expected on the basis of any one function
ll theories of concept learning assume that
ved and other types are lost or inaccessible;
has motivated an important component of
ve shall see, human conceptual behavior dis-
gory information (Medin & Ross, 1989), a
rototype theory.

y information abstracted from categories as
esentation discards information on category
lations among attributes (e.g., large spoons

a set of retrieval assumptions and an associated mapping onto performance. One could, for instance, assume that people store examples during learning but that new examples are classified by "computing" prototypes and comparing the similarity of the new example to these newly constructed prototypes. Indeed, it may be that the stored examples are so similar to each other that the *only* information people can successfully access is the central tendency or prototype. Such a model would have all the shortcomings of theories that assume that prototypes are formed during learning.

There are exemplar-based models of categorization that allow retrieval of more than the category central tendency (e.g., Hintzman, 1986; Medin & Schaffer, 1978), but they introduce a more serious problem for mixed models—they do not need prototypes to account for the data (e.g., Brooks, 1978, 1987). Exemplar models are perfectly capable of predicting that a newly presented prototype pattern may be classified more accurately than old examples (because it is highly similar to many same category members and has low similarity to members of contrasting categories) and that retention intervals will affect old patterns more than prototype patterns (e.g., Hintzman & Ludlum, 1980; Medin & Schaffer, 1978). Although these predictions may not be intuitively obvious, mathematical models based on these ideas provide both a qualitative and a quantitative account of these phenomena. This undercuts the initially strong underpinning for prototype theory. In head-to-head competition, exemplar models have been substantially more successful than prototype models (Busemeyer et al., 1984; Estes, 1986b, 1994; Medin & Smith, 1981; Nosofsky, 1988a,b, 1991, 1992; Shin & Nosofsky, 1992).

4. Summary

The initial flush of success of prototype theory proved to be a poor predictor of its future. One of the main functions of classification is that it allows us to make inferences and predictions on the basis of partial information. In general, the pairs of storage and retrieval assumptions associated with exemplar models preserve much more information than prototype models, information that people show sensitivity to. The context sensitivity of exemplar models is also consistent with much of the memory literature (e.g., Tulving, 1983).

The exemplar view of conceptual structure has a number of characteristics than distinguish it from other probabilistic view models. The prototype view claimed that categories were represented in terms of characteristic properties that worked together to create (linearly separable) categories where examples could be successfully classified on the basis of their similarity to prototypes. The exemplar view has no such requirement. The features used to categorize are the features of the category examples, and these need *not* be characteristic of the category overall. Some models with the Exemplar framework allow feature weighting to vary from example to example (e.g., Medin & Edelson, 1988). In short, the exemplar view appears to imply virtually no constraints on category membership. This is an issue we'll return to, but first we add a few more complications.

associated mapping onto performance. One
store examples during learning but that new
" prototypes and comparing the similarity of
ructed prototypes. Indeed, it may be that the
ther that the *only* information people can suc-
›r prototype. Such a model would have all the
hat prototypes are formed during learning.
›f categorization that allow retrieval of more
., Hintzman, 1986; Medin & Schaffer, 1978),
›blem for mixed models—they do not need
., Brooks, 1978, 1987). Exemplar models are
newly presented prototype pattern may be
ımples (because it is highly similar to many
similarity to members of contrasting cate-
affect old patterns more than prototype pat-
); Medin & Schaffer, 1978). Although these
vious, mathematical models based on these
quantitative account of these phenomena.
erpinning for prototype theory. In head-to-
ave been substantially more successful than
1984; Estes, 1986b, 1994; Medin & Smith,
›hin & Nosofsky, 1992).

: theory proved to be a poor predictor of its
assification is that it allows us to make infer-
›artial information. In general, the pairs of
ıated with exemplar models preserve much
›ls, information that people show sensitivity
models is also consistent with much of the

:ucture has a number of characteristics than
view models. The prototype view claimed
ns of characteristic properties that worked
ıtegories where examples could be success-
arity to prototypes. The exemplar view has
l to categorize are the features of the cate-
haracteristic of the category overall. Some
llow feature weighting to vary from exam-
1988). In short, the exemplar view appears
tegory membership. This is an issue we'll
omplications.

B. Challenges for Exemplar Approaches

We are nowhere near the end of our review. There's more to categorization than storing exemplars, and even if we restrict our attention to work with artificial categories, there have been developments that have served to motivate another basic shift in perspective. In this section we consider two of these developments.

1. Strategies and Rules

Neither prototype nor exemplar theories say anything about rules, but people who run as participants in laboratory studies of category learning often report that they are looking for and using rules. The structure of ill-defined categories is typically such that no simple rules are valid, but people may resort to more complex rules and strategies such as memorizing exceptions. There is also fairly good evidence that rules are not simply epiphenomena (e.g., Martin & Caramazza, 1980; Medin & Smith, 1981; Medin, Dewey, & Murphy, 1983; Medin, Wattenmaker, & Michalski, 1987; Nosofsky, Clark, & Shin, 1992) and models based on rules (e.g., Asby & Maddox, 1992; Nosofsky, Palmeri, & McKinley, 1994) may provide an excellent account of data involving artificial categories.

If people are using rules, why have exemplar models been so successful? There are three complementary possibilities that come to mind. One is that strategies do not eliminate exemplar coding, but they may bias it in a way that increases the parallels between rule-guided and exemplar-guided categorization (Medin, 1986; Medin & Wattenmaker, 1987). A second possibility is that the constraints associated with exemplar-based theories correspond to those of rule-based theories and the data reveal these common constraints (e.g., Wattenmaker, 1991, 1993). Third, rules and exemplar storage may be more or less independent, but performance may reflect a mixture of strategies (e.g., Brooks, 1987; Reagher & Brooks, 1993; Wattenmaker, 1993). Just how to model effects of these various strategies is a matter of considerable current concern (e.g., Estes, 1994; Kruschke, 1993; Nosofsky & Kruschke, 1992).

The influence of rules and strategies is reflected in other conceptual tasks as well. For example, Medin, Wattenmaker, and Hampson (1987) constructed artificial stimuli according to a family-resemblance principle and then asked participants to sort the stimuli into categories. The family-resemblance structure allowed people to create two categories, each organized around a prototype or best example (the resulting categories would be linearly separable). Over a range of stimuli and instructions Medin et al. failed to observe family-resemblance sorting. Instead, participants sorted on the basis of a single dimension. If the task structure prohibited unidimensional sorting, they would make an initial, partial sort on the basis of a single contrast and then employ a variety of strategies for dealing with the exceptions of leftovers. Sometimes these strategies will produce what appears to be family-resemblance sorting.

Follow-up studies showed that subjects could indeed be induced to produce family-resemblance sorting, and that a two-stage model embodying this general strategy (simple initial sort followed by dealing with leftovers) was able to predict when family-resemblance sorting would or would not be observed (Ahn & Medin, 1992; see also Spalding & Murphy, 1996; Reagher & Brooks, 1995). There are two implications of these observations. One is that participants do not simply assimilate probabilistic structures but rather organize them in terms of discrete structures plus noise. Second, a key characteristic of these discrete structures is that people prefer to create categories that are easy to describe. These observations underline the point that sorting is not simply driven by similarity computations. Whether or not this departure from similarity represents a bias associated with the communicative function of concepts, it represents a phenomenon not addressed by either prototype or exemplar models.

2. Feature Independence

The debate between prototype and exemplar theories has been about how featural information is integrated across examples and not about the nature of the features themselves. Indeed, predictions of category models depend crucially on being able to specify what the features are and having the research participants' analysis of stimuli agree with that of the experimenter. Researchers take care to ensure that agreement is realized either by using simple stimuli with salient dimensions (e.g., geometric shapes differing in size and color) or by employing multidimensional scaling techniques to identify (or confirm) dimensions (e.g., Nosofsky, 1987). Although this general strategy is absolutelv necessary for proper theoretical contrasts, it does carry with it the implicit assumption that the specific realization of abstract category structure is not important. To be sure, researchers are vigilant about establishing the generality of their results by employing a variety of stimuli. Nevertheless, this variety may tend to be biased toward realizations where the features of examples are independent and unrelated to each other (e.g., a triangle can be any color).

Are the properties or features of concepts generally independent? Probably not. First of all, the components of concepts are unlikely to have the status of being primitive features. Consider, for example, the features people typically list for the category *bird:* living, laying eggs, flying, having wings and feathers, singing, building nests, and so on. Each of these "features" is itself a complex concept with both an internal structure and an external structure based on interproperty relationships. For example, laying eggs implies a living organism and building nests is, in part, in service of protecting eggs. Flying requires wings and affords building nests in trees. In short, rather than independent features one has a web of relationships. This fact not only raises questions about how to define similarity over entities with interrelated features but also calls attention to issues of how the phenomena that we have been discussing might change as a function of nonindependent feature structure.

And change they do. Consider again, linear separability, which we mentioned could be characterized in terms of a summing of evidence against a criterion. In a series of experiments, Wattenmaker, Dewey, Murphy, and Medin (1986) found that linearly separable categories were easier to learn than nonlinearly separable cate-

ith leftovers) was able to predict when fam-
l not be observed (Ahn & Medin, 1992; see
& Brooks, 1995). There are two implications
ipants do not simply assimilate probabilistic
rms of discrete structures plus noise. Second,
tures is that people prefer to create categories
ions underline the point that sorting is not
s. Whether or not this departure from sim-
he communicative function of concepts, it
by either prototype or exemplar models.

nplar theories has been about how featural
es and not about the nature of the features
ory models depend crucially on being able
g the research participants' analysis of stim-
Researchers take care to ensure that agree-
stimuli with salient dimensions (e.g., geo-
or by employing multidimensional scaling
nsions (e.g., Nosofsky, 1987). Although this
r proper theoretical contrasts, it does carry
e specific realization of abstract category
earchers are vigilant about establishing the
variety of stimuli. Nevertheless, this vari-
ations where the features of examples are
(e.g., a triangle can be any color).
epts generally independent? Probably not.
s are unlikely to have the status of being
e, the features people typically list for the
having wings and feathers, singing, build-
res" is itself a complex concept with both
cture based on interproperty relationships.
organism and building nests is, in part, in
s wings and affords building nests in trees.
s one has a web of relationships. This fact
efine similarity over entities with interre-
sues of how the phenomena that we have
ion of nonindependent feature structure.
linear separability, which we mentioned
ming of evidence against a criterion. In a
ey, Murphy, and Medin (1986) found that
to learn than nonlinearly separable cate-

gories, if the stimuli or instructions facilitated interproperty coding that was compatible with a summing of evidence (features). An example from one of their studies involved a category whose modal properties were "made of metal," "medium-sized," "has a regular surface," and "easy to grasp." The contrasting category had differing modal values on each of these dimensions. Out of context, one can think of many interproperty relationships, and no one appears to be particularly salient. In category-learning conditions run without further instructional elaboration, no advantage for linearly separable categories was observed. In a second condition, however, participants were given the further hint that the objects in one category might serve as a substitute for a hammer. This hint improved performance overall to some extent, but most striking was the observation that the linearly separable category structure was now much easier than the nonlinearly separable structure. This result might appear to support prototype theory, but remember that prototype theory requires featural independence (nor can it explain when linear separability matters and when it does not).

This selective facilitation of linearly separable categories reflects something beyond making the stimulus materials more meaningful. In other studies Wattenmaker et al. showed that instructional hints could selectively improve performance on nonlinearly separable categories—the key factor was the relationship between the types of interproperty encoding induced and the category structure. A clear implication of these findings is that one cannot make generalizations in terms of abstract category structures and expect them to carry through across contexts, because different types of interproperty or relational coding may take place (see Medin et al., 1987, and Murphy & Spalding, 1996, for corresponding observations from sorting tasks. Family-resemblance sorting is readily observed under appropriate relational coding conditions). Equally significant is the point that interproperty relationships are outside the boundary conditions of almost all current categorization models (including prototype, exemplar, and rule-based models). Therefore, these models currently have limited generality, and this limitation is most evident where one might most want to generalize—meaningful stimuli.

In summary, although exemplar models show a lot of promise (see Smith & Zarate, 1992, for a recent informative extension of exemplar models to social categories), this section has revealed two important problems and limitations. One is that people employ rules and strategies, even for artificial, relatively meaningless stimuli. The other limitation (which also applies to all the other models that have been under discussion) is that exemplar models have only been developed for contexts where the features of category members are independent. That is, interproperty relationships are not addressed. The latter problem is quite serious, as one would like to move freely between natural and artificial stimuli in our analyses of concepts. To do this, we need to address meaningful stimuli and attendant issues of relational coding. In the next section this issue will be developed further, both with respect to ideas about similarity and with respect to the role of knowledge structures in category organization.

C. General Constraints: Similarity and Theories

A series of related concerns, none by itself perhaps too serious, has culminated in a fundamental shift away from feature-based similarity models. This has given rise to two major trends that appear to be contradictory but actually may serve to reinforce each other. One important position is that similarity is too unconstrained to perform a useful explanatory function and that, instead, the claim is that categories are organized in terms of theories about the world (e.g., Murphy & Medin, 1985). The alternative perspective argues that we need better theories of similarity (e.g., Goldstone, 1994a, 1995; Medin, Goldstone & Gentner, 1993).

1. Similarity

Why do we have the categories we have and not others? A major problem with using similarity to explain categories is that similarity is too flexible. For example, in Tversky's (1977) contrast model similarity is a weighted function of shared and distinctive features. Similarity will depend critically on the weights given to particular features or properties. To borrow an example from Medin (1989), a *zebra* and a *barber pole* would be more similar than a *zebra* and a *horse* if the feature "striped" is weighted sufficiently. To complicate the picture, Tversky and others (e.g., Gati & Tversky, 1984; Landau, Smith, & Jones, 1988; Medin, Goldstone, & Markman, 1995, for a review) have demonstrated that the weight given to a feature depends on the context instructions, experimental task, and even the concept under consideration (Ortony, Vondruska, Foss, & Jones, 1985).

What's wrong with flexibility? The key issue concerns explanatory power. To the extent that similarity is an outcome of a series of other processes, it is more a dependent variable than an independent variable. That is, the processes that determine feature weighting are doing the explanatory work, and the general appeal to similarity only serves to conceal that fact.

A related set of problems derives from determining what counts as a feature. Returning to the Rosch and Mervis (1975) studies, it would be a mistake to assume that people asked to list features of concepts were able to "read" and report their mental representations in an accurate manner. Indeed Keil (1979, 1981) noted that category examples (e.g., *robin, ostrich*) share many important properties that virtually never appear in feature listing (e.g., has a heart, has blood, sleeps, occupies space, can be thought about, etc.). Keil argued that knowledge about just these sorts of predicates serves to organize children's conceptual and semantic development. In short, to understand feature listings we need a process model for access to conceptual knowledge.

To take things a step further, one could argue that without constraints on what is to count as a feature, any two things can be arbitrarily similar or dissimilar. Thus as Murphy and Medin (1985) suggest, the number of features that *plums* and *lawnmowers* share could be unlimited: both are found on Earth, both are found in our solar system, both are touched by people, both do not see or hear well, both do not

and Theories

tself perhaps too serious, has culminated in
·based similarity models. This has given rise
ontradictory but actually may serve to rein-
.on is that similarity is too unconstrained to
and that, instead, the claim is that categories
it the world (e.g., Murphy & Medin, 1985).
we need better theories of similarity (e.g.,
tone & Gentner, 1993).

ave and not others? A major problem with
; that similarity is too flexible. For example,
ilarity is a weighted function of shared and
nd critically on the weights given to partic-
an example from Medin (1989), a *zebra* and
n a *zebra* and a *horse* if the feature "striped"
the picture, Tversky and others (e.g., Gati &
es, 1988; Medin, Goldstone, & Markman,
that the weight given to a feature depends
ntal task, and even the concept under con-
Jones, 1985).
cey issue concerns explanatory power. To the
series of other processes, it is more a depen-
iable. That is, the processes that determine
atory work, and the general appeal to simi-

rom determining what counts as a feature.
975) studies, it would be a mistake to assume
ncepts were able to "read" and report their
nanner. Indeed Keil (1979, 1981) noted that
hare many important properties that virtu-
has a heart, has blood, sleeps, occupies space,
ed that knowledge about just these sorts of
conceptual and semantic development. In
need a process model for access to concep-

ould argue that without constraints on what
can be arbitrarily similar or dissimilar. Thus
the number of features that *plums* and *lawn-*
are found on Earth, both are found in our
le, both do not see or hear well, both do not

shave, both can be dropped, and so on (see also Goodman, 1972; Watanabe, 1969). The general point is that attempts to describe conceptual behavior in terms of similarity will prove useful only to the extent that one specifies what is to count as a relevant feature and which principles determine the importance of particular properties. And the explanatory work is being done by the principles that specify these constraints, not some amorphous notion of similarity.

Even when similarity can be nailed down, it may not explain conceptual behavior. Consider, for example, experiments by Gelman and Markman (1986), which pitted category membership against perceptual similarity in an inductive reasoning task. Children were first shown pictures of two animals and taught that different (novel) properties were true of them. They were then asked which property was true of a pictured new example where the example was perceptually similar to one of the first pictures but shared category memberships with the (less similar) other example. Children judged that the new example would have the property of the animal that was of the same category but perceptually different. It appears that similarity can be overridden by other forms of knowledge.

Studies with adults have identified conditions under which similarity is neither necessary nor sufficient to determine category membership (e.g. Rips, 1989; Rips & Collins, 1993). One could argue that the problem is that similarity is so dynamic that similarity in a similarity rating task is different from similarity in a categorization task. Although this argument could in principle have validity, in practice it tends to undermine using similarity to understand conceptual behavior. That is, the burden of proof shifts to similarity theorists to develop computational models that address these conjectured interactions of similarity with tasks.

2. Theories

Many philosophers of science have argued that observations are necessarily theory laden, and recently researchers have begun to adopt the framework that conceptual behavior may be knowledge-based and driven by theories (e.g. Carey, 1985; Keil, 1986; Lakoff, 1987; Markman, 1987; Massey & Gelman, 1988; Medin, 1989; Murphy & Medin, 1985; Oden, 1987; Rips, 1989; Schank, Collins, & Hunters, 1986; Wisniewski & Medin, 1991, 1994; and see Komatsu, 1992 and Medin & Heit, in press, for reviews, and Hirschfeld & Gelman, 1994; VanMechlen, Hampton, Michalski, & Theuns, 1993, and Nakamura, Taraban, & Medin, 1993 for relevant edited volumes). One difference in perspective is that according to the theory-based view, categorization is not simply based on a direct matching of properties of a concept with those of an example, but rather requires that the example have the right "explanatory relationship" to the theory organizing the concept. Informally one could say that the relationship between a concept and an example is like the relationship between theory and data. The general notion is that many concepts may be organized around features that are more abstract than the (perceptible) features of examples. Information from examples may be used to infer the presence of more abstract, possibly causal properties. For example, there is no simple checklist of properties that define a person as an introvert or extrovert; features that typically are linked to these concepts (staying home on a Friday night implies introversion) may be "blocked" or defeated by other information (i.e., the person may be recovering from a wild night of partying on Thursday).

Theories help determine which properties are relevant to a categorization task. One way of reconciling similarity and explanation is to argue that similarity operates on features selected by theories. The nice feature of this view is that laboratory studies with artificial knowledge-poor stimuli could be informative with respect to how conceptual behavior proceeds once the relevant set of features has been determined. Wisniewski and Medin (1994) argue, however, that more tightly coupled systems are needed. Theories may do more than act as filters determining feature relevance. Wisniewski and Medin gave participants categorization and rule induction tasks where the examples (children's drawings of people) were associated with different domain theories cued in by different category labels. For example, in one condition participants were told that the drawings were done by creative versus noncreative children but in another participants were told that these same drawings were done by mentally healthy versus emotionally disturbed children. The results show, first of all, that it is not reasonable to assume some *a priori* set of unambiguous features. The "features" comprising participant's category descriptions varied as a function of category labels and the same aspect of drawing was interpreted differently for different labels. In addition, participants sometimes reinterpret features when given feedback about category membership. For example, they might have classified a drawing as one by a creative child because it did not have the usual simple smile, but when told that the drawing was done by a noncreative child the participant might decide that how the mouth was drawn reveals that the child was unhappy about their lack of creativity. Finally, participants' rules often involved abstract features that are operationalized differently as a function of learning history (e.g., how detailed does a drawing have to be to qualify as "detailed?"). In the absence of meaningful category labels, the rules of participants involved fairly superficial features in the drawings (e.g., buttons visible on shirt); when the labels were meaningful, participants appeared to treat the labels as implying causes where properties of drawing were "effects." In the latter case learning consisted of establishing linkages between ideas about the categories and properties of the children's drawings. On the basis of these observations, Wisniewski and Medin argued that theories not only select features but also "create" features and determine how they will be instantiated. This is consistent with the general view that categorization is not simply a syntactic matching of features between concepts and examples.

A key aspect of the theory-based view is that it may address the question of why we have the categories we have. Coherence may even be achieved in the absence of any obvious source of similarity among category examples. The category composed of money, pets, photo albums, and children seems unusual to say the least. In the context of a goal-derived category of "things to take out of one's house in case

:y

xample, there is no simple checklist of prop-
ert or extrovert; features that typically are
on a Friday night implies introversion) may
mation (i.e., the person may be recovering
day).
perties are relevant to a categorization task.
xplanation is to argue that similarity oper-
e nice feature of this view is that laboratory
timuli could be informative with respect to
the relevant set of features has been deter-
argue, however, that more tightly coupled
nore than act as filters determining feature
participants categorization and rule induc-
's drawings of people) were associated with
fferent category labels. For example, in one
he drawings were done by creative versus
icipants were told that these same drawings
emotionally disturbed children. The results
e to assume some *a priori* set of unambigu-
participant's category descriptions varied as
me aspect of drawing was interpreted diff-
participants sometimes reinterpret features
nembership. For example, they might have
child because it did not have the usual sim-
ig was done by a noncreative child the par-
outh was drawn reveals that the child was
. Finally, participants' rules often involved
differently as a function of learning history
ve to be to qualify as "detailed?"). In the
ne rules of participants involved fairly super-
ons visible on shirt); when the labels were
at the labels as implying causes where prop-
atter case learning consisted of establishing
ries and properties of the children's draw-
, Wisniewski and Medin argued that theo-
eate" features and determine how they will
the general view that categorization is not
between concepts and examples.
w is that it may address the question of why
ence may even be achieved in the absence
ng category examples. The category com-
d children seems unusual to say the least. In
of "things to take out of one's house in case

of a fire," the category immediately becomes sensible (Barsalou, 1983; see also Barsalou, 1985) and could even be projected to predict further instances (valuable and personal papers, furniture passed down from relatives, etc.).

There is evidence that even very young children's category learning is constrained by domain-specific theoretical biases (e.g., Keil, 1989). Carey (1985) concurs and suggests that young children's biological reasoning is organized by a form of naive psychology. She has evidence that 5–7 year old's inferences are guided by similarity to humans, not an unbiased overall similarity. For instance, they are more sure that a *bug* has some novel property if they are told that *people* have it than if they are told *bees* have it. More recent work (Coley, 1995; Inagaki & Hatano, 1993) challenges Carey's claim about the psychological nature of children's early explanatory theory about living things. However, note that the debate is over the specific nature of the theory; it is agreed that explanatory theories play a crucial role in organizing the concepts of even preschool children.

It may be that people approach all categories with something of a theoretical stance in the sense of having generalized expectations concerning the basis for category membership. For example, people might believe that the key organizing principle for artifacts is function or the intentions of the builder but view biological kinds as depending on intrinsic underlying properties (e.g., Atran, 1990; Keil, 1989; Rips, 1989; Malt, 1990, 1994; Medin & Ortony, 1989; Gelman & Wellman, 1991). In the case of biological kinds, people appear to believe that a true underlying nature or essence that imparts category identity. This view has been dubbed *psychological essentialism* because it is concerned with people's assumptions about the world, not how the world truly is (with respect to the latter, see Mayr, 1988, for a review of how essentialist biases have slowed the development of modern biological thought).

What is the basis for these claims about psychological essentialism? Gelman, Coley, and Gottfried (1994) point to four types of evidence as relevant: (a) appeal to invisible causal mechanisms to explain appearance and changes associated with growth, (b) the assumption of innate dispositions or inborn capacities to explain capacities that emerge later, (c) belief in the maintenance of identity despite changes in superficial appearance, and (d) the assumption that members of a category share a large number of other properties (i.e., that these categories have rich induction potential). Gelman et al. (1994) review supportive evidence for each of these assumptions (see also Shipley, 1993).

Overall, then, there is considerable evidence that knowledge, theories, and belief systems affect categorization and reasoning in a manner unanticipated by the classical, probabilistic, or exemplar views of categorization. In some ways these views seem orthogonal to the theory-based perspective because they seem to be about structure, whereas theories seem to be more about content. We are not so sure such a glib distinction can or should be made because content may determine how structure is encoded and instantiated. What does seem clear is that the notion of independent features embodied in all three views may be off the mark for many (if not most) categories. The question we take up in the next section is whether the cumulative evidence suggests that the marriages between similarity models and theories of conceptual behavior and between studies with artificial categories and studies with natural categories should be broken up.

3. Implications

What do recent views on theory-based categories mean for research in artificial categories? Are findings from studies using random dot patterns or large, red triangles as stimuli relevant to the rich causal structure that links the observation that *robins* have hollow bones to the expectation that *eagles* have hollow bones? It is tempting to argue that stimuli associated with laboratory studies of artificial categories are simply too impoverished and that theories based on such studies will likewise be impoverished. Not only are the stimuli impoverished, but they have typically been constructed to have independent, unrelated features. Strike two.

But not strike three. We believe that progress is cumulative but difficult, and we should not abandon any tool without careful consideration. A shift in theoretical orientation does not nullify previous research or the potential insights growing out of it. Consider, for example, the shift from the classical to the probabilistic view. Laboratory studies of rule learning associated with well-defined categories seemed, on the surface, irrelevant to ill-defined categories where any rules would have numerous exceptions. Nonetheless, when researchers began to pay attention to the strategies learners employ for probabilistic categories, rules again received attention. Furthermore, even when a theoretical orientation does not survive intact, some of its insights may be passed on. For example, early work on hypothesis testing carried the key idea of selective attention, an assumption that now is taken for granted and embodied in virtually all current categorization theories.

A pragmatic reason to avoid a divorce between artificial and natural categories is that there is no simple line that can be drawn between them. Is a large red triangle artificial and a children's drawing of a person natural, or does the drawing become a natural stimulus only when it is associated with a meaningful category label? In addition, just as researchers found that research participants imported knowledge to render nonsense syllables meaningful, so also do learners bring knowledge and bias to organize random dot patterns (Hock, Webb, & Cavedo, 1987). It seems to us that there is a continuum of research strategies available that may trade theoretical precision and experimental control for richness and generalizability. In short, we would be very much surprised if laboratory studies with so-called artificial stimuli did not continue to inform the discussion of concepts and categories.

At the same time, however, we do believe that the fruitfulness of work with artificial categories (that is, work tending toward the artificial end of the artificial–natural continuum) is directly tied to our understanding of natural categories. What we've learned from research on natural categories could be usefully applied to research with artificial categories. Researchers are far more likely to use ill-defined

he marriages between similarity models and
:tween studies with artificial categories and
›e broken up.

:ategories mean for research in artificial cat-
; random dot patterns or large, red triangles
ıcture that links the observation that *robins*
ıat *eagles* have hollow bones? It is tempting
boratory studies of artificial categories are
ries based on such studies will likewise be
impoverished, but they have typically been
ıted features. Strike two.
progress is cumulative but difficult, and we
areful consideration. A shift in theoretical
earch or the potential insights growing out
om the classical to the probabilistic view.
iated with well-defined categories seemed,
categories where any rules would have
n researchers began to pay attention to the
c categories, rules again received attention.
ientation does not survive intact, some of
e, early work on hypothesis testing carried
ımption that now is taken for granted and
ization theories.
between artificial and natural categories
drawn between them. Is a large red trian-
f a person natural, or does the drawing
is associated with a meaningful category
und that research participants imported
meaningful, so also do learners bring
ı dot patterns (Hock, Webb, & Cavedo,
nuum of research strategies available that
imental control for richness and general-
ı surprised if laboratory studies with so-
o inform the discussion of concepts and

ve that the fruitfulness of work with arti-
ırd the artificial end of the artificial–nat-
:derstanding of natural categories. What
categories could be usefully applied to
ıers are far more likely to use ill-defined

than well-defined categories because they believe that the classical view is less viable than the probabilistic view. Likewise, researchers might profitably come to favor stimuli with nonindependent, structured constituents rather than restricting themselves to the narrow set of situations where the assumption of independent features may be viable. In a similar vein, the categorization function of concepts is but one of many, yet work with artificial categories has focused almost exclusively on it. Although it may take a bit of ingenuity, we see no principled reason why virtually every conceptual function could not be instantiated with (relatively) artificial stimuli. Finally, we think that laboratory studies favoring artificial categories and stimuli (again, employing some ingenuity) could create more effective parallels between conditions in the laboratory and outside of it. This form of scrutiny has been useful in the past, and we see no reason why it could not be effective. In short, we support a continuing dialogue and interaction between the artificial and the natural.

What about the implications of theory-based categories for models of similarity? To be honest, we see only a limited future for similarity-based models based on relatively unconstrained, independent features. For purposes of some analyses one only needs to be in the right ballpark, and feature-based models are easy to work with. More generally, however, objects and events of interest are composed of more than a list of properties—instead, they are structured in terms of a variety of interproperty and even hierarchical relationships. The evidence is mounting that we need more powerful models of similarity that address not only structure but also processing principles (e.g., Goldstone, Medin, & Gentner, 1991; Gentner & Markman, 1995; Goldstone, 1994b; Goldstone & Medin, 1994; Holyoak & Thagard, 1991; Markman & Gentner, 1993a,b; Gentner & Markman, 1994; Medin et al., 1993).

Will enriched models of similarity be able to address the phenomena associated with research on contextually rich theory-driven categorization considered earlier? Our opinion is that dynamic, structural models are at least necessary; whether they are also sufficient remains to be seen (for more positive readings see Goldstone, 1994a; Jones & Smith, 1993, but also Gelman & Medin, 1993). It also remains to be seen whether similarity-based models will constitute an important component of conceptual behavior or whether their key ideas will be so closely integrated with theory-based considerations that one will not be able to isolate any one module and label it as "similarity." In any event, far from abandoning models of similarity, we would argue that considerably more energy should be directed at their further development, including their more context-dependent creative aspects.

D. Summary

We have had to strain chronology quite a bit to force a number of uneven, parallel developments into the story as so far presented. And much has been left out (for reviews from a different perspective see Goldstone, 1994a,b; Homa, 1984; Komatsu, 1992; Medin & Heit, in press; Oden, 1987; Rips, 1990). Still, we do not think we have strayed too far in pointing to the shifts from the classical to the probabilistic view, and the ensuing tensions between prototype and exemplar models and between similarity-based and knowledge-driven categorization. Missing so far, however, has been an account of what we might call a form of "adaptive radiation" of work on conceptual behavior. Not only are a broader range of conceptual functions being studied, but researchers are also focusing on new and distinct domains of inquiry. This focus on domains represents less a lack of faith in general principles of categorization than a positive conviction that certain phenomena and principles may be more readily seen in some contexts than others. In any event, we see adaptive radiation as a positive development so long as there is a flow of ideas and information across psychological niches. The final main section of this chapter briefly samples some of this ongoing activity.

VII. CONCEPTUAL FUNCTIONS AND COGNITION IN CONTEXT

In this final section we have space to do little more than point to some current salutary trends. Questions about concepts and categorization are central to understanding the nature of mind. The burgeoning current activity may be distressing to those who think that a monolithic enterprise is a guarantor of coherence and quality, but we see the diversity as an unequivocal healthy sign.

A. Perception and Representation

Readers of the present volume may have been asking themselves the question of what the difference is between object recognition and categorization. It seems obvious that object recognition *is* a form of categorization and that, therefore, there should be an ongoing interchange of ideas between object recognition and categorization research. Although there has been far less interplay than one might imagine, there are some encouraging signs.

First of all, there is at least some convergence in terms of theoretical ideas. The influential geon theory of object recognition (Biederman, 1985; Hummel & Biederman, 1992) aims to recognize objects at the basic level and incorporates both components (geons) and relations between them. Template models, an alternative to the geon approach, incorporates structure implicitly (e.g., Tarr, 1995). Most recently, researchers have begun to apply ideas and models from object recognition to phenomena such as typicality judgments and categorization reaction times (Kurbat, 1995). There is even evidence that measures of shape similarity predict unique variance in situations where the names of objects rather than the objects themselves are presented (Kurbat, 1995). This suggests that perceptual representations may be activated and used for verbal categorization and reasoning (see also Barsalou, 1993; Barsalou, Solomon, & Wu, in press).

Ideas deriving from categorization are also relevant to object recognition. For example, there is increasing evidence that categorization is not a passive consumer of features or components but instead directly influences feature construction

ween prototype and exemplar models and
edge-driven categorization. Missing so far,
we might call a form of "adaptive radiation"
only are a broader range of conceptual func-
e also focusing on new and distinct domains
resents less a lack of faith in general princi-
onviction that certain phenomena and prin-
ne contexts than others. In any event, we see
oment so long as there is a flow of ideas and
nes. The final main section of this chapter
ctivity.

AND COGNITION IN CONTEXT

little more than point to some current salu-
s and categorization are central to under-
eoning current activity may be distressing to
erprise is a guarantor of coherence and qual-
uivocal healthy sign.

ave been asking themselves the question of
cognition and categorization. It seems obvi-
of categorization and that, therefore, there
ideas between object recognition and cate-
been far less interplay than one might imag-

vergence in terms of theoretical ideas. The
nition (Biederman, 1985; Hummel & Bie-
cts at the basic level and incorporates both
veen them. Template models, an alternative
ructure implicitly (e.g., Tarr, 1995). Most
y ideas and models from object recognition
nts and categorization reaction times (Kur-
measures of shape similarity predict unique
of objects rather than the objects themselves
ests that perceptual representations may be
ion and reasoning (see also Barsalou, 1993;

re also relevant to object recognition. For
at categorization is not a passive consumer
d directly influences feature construction

(Goldstone, 1995c; Schyns & Murphy, 1994; Thibaut & Schyns, 1995; Wisniewski & Medin, 1994). That is, people may not come with a present vocabulary of features but rather features may be learned. If so, then theories of object recognition need to be sensitive to learning history and the possibility that the nature and importance of features might change across domains (cf. Hochberg, chap. 9, this volume).

Overall, there has been limited contact across areas but what there is, is very promising. We predict increasing interaction between mainstream perception and mainstream categorization (see also Harnad, 1987).

B. Conceptual Functions beyond Categorization

Another promising avenue of research is more in-depth exploration of conceptual functions beyond categorization. We will look at two such functions—conceptual combination and induction—in a little more detail. For work on other conceptual functions see Barsalou (1993); Markman, Yamauchi, and Makin (1997); Murphy and Wisniewski (1989); and Wisniewski, (1995).

1. Conceptual Combination

Conceptual combination is important because it allows a productive use of conception. From a vocabulary of simple concepts we can create and understand a potentially unlimited set of combined concepts. Although we are rarely faced with the task of understanding a single concept in isolation, most categorization models do not address combined concepts (Rips, 1995). Conceptual combination is also important because it requires an analysis of conceptual structure and of relations between concepts (e.g., Osherson & Smith, 1981; Smith & Osherson, 1984). For example, a *blackboard* is judged to be a member of the combined category *school furniture* but not a member of the simple category *furniture* (Hampton, 1982, 1987, 1988).

It is far from easy to develop process models for conceptual combination. Nonetheless, a clear, but limited model may point to where progress is needed. Consider the selective modification model of Smith, Osherson, Rips, and Keane (1988), which aims to address adjective noun combinations such as *green apple.* According to this model, to understand this combination a person would retrieve the prototype representation for *apple,* pay extra attention to the dimension of color, and replace the default value of red for applies with the value green. In brief, the person would be constructing a new prototype for *green apple* by modifying the apple prototype. The selective modification model then uses the *green apple* prototype and associated weightings to predict categorization judgments and typicality ratings of potential examples of the new concept.

Although the modification model has enjoyed some success, conceptual combination is more complex than the model allows. People use their general knowledge about relations among features so that dimensions other than the one named by the adjective may be affected. For example, a *brown apple* is not only of an atypical color, but people may infer that it is rotting (Medin & Shoben, 1988). Indeed, conceptual combination sometimes leads to inferences about emergent features that are not expected to be true of either constituent concept taken singly (Hampton, 1987; Hastie, Schroeder, & Weber, 1990; Kunda, Miller, & Claire, 1990; Murphy, 1988; Rips, 1995). For example, people expect neither a *carpenter* nor a *Harvard-educated person* to be a nonconformist, but they do expect a *Harvard-educated carpenter* to be a nonconformist (Kunda et al., 1990). These observations suggest that knowledge outside of the two constituent concepts is brought to bear in understanding combinations.

If understanding adjective–noun combinations is hard, noun–noun combinations are even harder. One interesting approach to noun–noun combinations is the idea that people align the nouns, then attribute a salient property of the modifying noun to the head noun (Wisniewski & Gentner, 1991). Thus a *skunk squirrel* may be a smelly squirrel and a *zebra horse* may be interpreted as a horse with stripes. Another strategy that Wisniewski (1996) calls "relation-linking" seems to emerge when the constituent concepts are less similar. For example, a *skunk box* may be interpreted as a box for containing skunks rather than as a smelly box. The idea of analyzing conceptual combination into constituent strategies is appealing, but the challenge is to then integrate these processes into a unitary model (see Wisniewski, 1997, for some ideas along these lines). Overall, conceptual combination is emerging as a challenging and important area of research.

2. Induction

Concepts are used in reasoning. The issue of how we infer novel properties of categories has been addressed by Osherson, Smith, and their associates in the form of a category-based induction model (CBIM). This model was inspired by earlier work on category-based induction by Rips (1975). For example, suppose that you are told that *cows* have some enzyme in their blood and you are asked to assess how likely it is that *bears* also have this property. According to the CBIM, two factors determine the soundness of such inductive inferences. The greater the similarity between the premise category (e.g., *cow*) and the conclusion category (e.g., *bears*) the more confident one should be that the inference carries over.

The second factor is the *coverage* of the premise that is defined as the similarity between the category (or categories) in the premise and members of the lowest level superordinate category that encompasses the categories in the premise and conclusion. The most specific category that includes cows and bears is *mammal*. The category *cow* is fairly similar to other members of the category *mammal* (cows are considered to be typical mammals). Thus if cows have some enzyme in their blood, it is plausible that all mammals might have this enzyme and bears are mammals. According to the CBIM, the overall confidence in this induction is a weighted average of the similarity and coverage components.

·, a *brown apple* is not only of an atypical color, 'Medin & Shoben, 1988). Indeed, conceptual rences about emergent features that are not uent concept taken singly (Hampton, 1987; ında, Miller, & Claire, 1990; Murphy, 1988; ›ect neither a *carpenter* nor a *Harvard-educated* y do expect a *Harvard-educated carpenter* to be . These observations suggest that knowledge ›ts is brought to bear in understanding com-

:ombinations is hard, noun–noun combina- approach to noun–noun combinations is the attribute a salient property of the modifying & Gentner, 1991). Thus a *skunk squirrel* may may be interpreted as a horse with stripes. 96) calls "relation-linking" seems to emerge ss similar. For example, a *skunk box* may be ınks rather than as a smelly box. The idea of o constituent strategies is appealing, but the ›cesses into a unitary model (see Wisniewski, . Overall, conceptual combination is emerg- ·a of research.

sue of how we infer novel properties of cat- ›n, Smith, and their associates in the form of ·IM). This model was inspired by earlier work 1975). For example, suppose that you are told lood and you are asked to assess how likely it :ording to the CBIM, two factors determine ·nces. The greater the similarity between the nclusion category (e.g., *bears*) the more con- ɛ carries over.

` the premise that is defined as the similarity . the premise and members of the lowest level ;es the categories in the premise and conclu- ncludes cows and bears is *mammal*. The cat- ıbers of the category *mammal* (cows are con- if cows have some enzyme in their blood, it have this enzyme and bears are mammals. nfidence in this induction is a weighted aver- ıponents.

One might wonder why the coverage component of the model is even needed. Isn't similarity enough? The answer is that coverage accounts for some phenomena that the notion of similarity by itself may not. For example, inductive inferences are stronger when they go from a typical category member (e.g., for mammals, *cow*) to an atypical member (e.g., *anteater*) than in the reverse direction (Rips, 1975). The CBIM assumes that similarity is symmetrical (*b* is as similar to *a* as *a* is to *b*) but the coverage component accounts for the asymmetry in reasoning. The category *cow* is more similar to other members of *mammal* than is the category *anteater;* therefore, inferences from *cow* are stronger than inferences from *anteater.*

The CBIM accounts for a variety of phenomena involving how people use similarity and category information in making inferences (e.g., Smith, Shafir, & Osherson, 1993). So far, however, the model has focused on the role of categories and has not addressed the role of properties. Specifically, the model is tested using "blank predicates," (e.g., "requires biotin for hemoglobin synthesis"). These properties are chosen so that subjects are likely to have few prior beliefs about how the properties are linked to specific categories. It seems clear, however, that even at an abstract level, "kinds" of properties may influence induction. The "blank predicates" used by Osherson et al. were almost all physiological in nature. That may have been enough information to lead subjects to highlight certain relations at the expense of others. For example, Heit and Rubinstein (1994) found that different similarity relations guided inferences about unfamiliar behavioral properties versus information about unfamiliar anatomical properties. Thus, patterns of induction probably depend crucially on the content of the properties. "Blank" properties may not be truly blank, and an important challenge for models of induction is to incorporate the information embodied in meaningful properties (for some initial ideas along these lines, see Smith et al., 1993). Overall, the CBIM and associated studies have sparked a resurgence of interest in induction.

C. Language, Categories, and Induction

Another promising trend we would like to highlight is the careful examination of the impact of language on conceptual functions. Most of this research has been done with children, but results reveal that language has a powerful organizing effect on conceptual structure and function. Specifically, we will discuss research indicating an early impact of language on categorization and induction.

1. Categorization

One line of work bearing out this point has been done by Waxman and colleagues (Waxman, 1990; Waxman & Gelman, 1986; Waxman & Markow, 1995). Waxman and colleagues show that children use linguistic cues to establish conceptual hierarchies, and that the effects of language vary depending on the hierarchical level and the perceptual support for the category being learned. Specifically, providing a label for a category helps children to sort at the superordinate level. For instance, children asked to sort objects into superordinate classes (e.g., *food, animals*) did better when given a novel (Japanese) label for that class than when given no label at all (Waxman & Gelman, 1986). Moreover, Waxman (1990) found that while providing a novel label for a category (e.g., *dobu*) enhanced preschoolers' sorting at the superordinate level, it had little effect on sortings at the basic level, and actually interfered with sorting at the subordinate level. In contrast, labeling a category with a novel adjectival phrase (e.g., *dobish ones*) enhanced sorting at subordinate levels, but interfered with sorting at superordinate levels.

In summary, nouns organize higher-order categories, and appear to emphasize the commonalities among disparate objects; adjectives organize lower-order categories and emphasize differences between similar objects. Interestingly, performance for basic-level categories tested by Waxman and colleagues is consistently at ceiling; apparently, salient perceptual cues are sufficient to induce classification at the basic level without additional support. These results suggest that the conceptual organizing role played by language is specific as well as powerful.

D. Expertise and Culture

Another salutary trend is a growing focus on how knowledge impacts conceptual systems. This takes the form of examining both how expertise effects reasoning in a particular domain, and investigating conceptual similarities and differences among members of different cultures.

1. Expertise

Psychologists have become increasingly interested in studying expertise (see Bedard & Chi, 1992, for a review). One way that expertise is useful is in comparing novices to experts. Studies of expertise have revealed profound reorganizations of knowledge as a function of amount and type of expertise (e.g., Chi, 1992; Chi & Bjork, 1991; Medin, Lynch, Coley, & Atran, 1997; Tanaka & Taylor, 1991). Such studies potentially reveal both how knowledge is structured in a well-known domain, and what conceptual changes lead to that structure. We feel this is an important and understudied aspect of conceptual functioning. Most of our contact with the world may involve domains with which we are relatively familiar. Indeed, we find it conceivable that further investigations of expertise will reveal that some if not many of the findings of cognitive psychology are more accurately seen as novice heuristics than as general characterizations of conceptual functioning.

Another way in which expert subjects can inform theories of conceptual organization is by allowing the interplay of mind and world in shaping conceptual structure to be examined. By comparing experts from different subfields within the same domain, the subject matter ("world") can be held constant, but the knowledge, goals, and naive theories of the experts ("mind") may well vary. In this way the con-

the superordinate level. For instance, chil-
dinate classes (e.g., *food, animals*) did better
r that class than when given no label at all
, Waxman (1990) found that while provid-
lobu) enhanced preschoolers' sorting at the
sortings at the basic level, and actually inter-
vel. In contrast, labeling a category with a
enhanced sorting at subordinate levels, but
levels.

-order categories, and appear to emphasize
jects; adjectives organize lower-order cate-
n similar objects. Interestingly, performance
nan and colleagues is consistently at ceiling;
fficient to induce classification at the basic
results suggest that the conceptual organiz-
s well as powerful.

us on how knowledge impacts conceptual
ng both how expertise effects reasoning in
nceptual similarities and differences among

nterested in studying expertise (see Bedard
expertise is useful is in comparing novices
ealed profound reorganizations of knowl-
of expertise (e.g., Chi, 1992; Chi & Bjork,
997; Tanaka & Taylor, 1991). Such studies
s structured in a well-known domain, and
ructure. We feel this is an important and
oning. Most of our contact with the world
relatively familiar. Indeed, we find it con-
pertise will reveal that some if not many of
more accurately seen as novice heuristics
eptual functioning.

s can inform theories of conceptual orga-
ind and world in shaping conceptual struc-
rts from different subfields within the same
an be held constant, but the knowledge,
mind") may well vary. In this way the con-

tributions of mind and world may be assessed. Along these lines, we are currently undertaking a large-scale study of tree expertise (Medin et al., 1997). We examine how landscapers, maintenance workers, and taxonomists categorize and reason about the same set of trees. Therefore, any differences we find cannot be attributed to difference in the content domain and must instead be attributed to differences in the conceptual structure of our experts. Care must still be exercised; similarities between kinds of experts are not necessarily due to common perceptions of structure in the world. For instance, our experts agree on clusters for a subset of the tree species we used as stimuli, but differ systematically on their justifications for these clusters. Almost all experts grouped black walnut with shagbark hickory, but based on justifications, landscapers did so because both are large, attractive trees useful as a centerpiece in a landscape layout, maintenance workers did so because both are nut trees, and taxonomists did so because both belong to the scientific family *Juglandaceae.* True, these facts about black walnut and shagbark hickory are not unrelated, but apparently different components of the correlational structure of the world was salient to different kinds of tree experts. Our studies also show a role of language in reasoning, at least among one of the subgroups. Landscapers and taxonomists privilege the genus level in reasoning but maintenance workers do so only when the common term marks the shared genus (e.g., white oak and burr oak are both oaks, belonging to the genus, *Quercus* versus pairs like eastern cottonwood and white popular, which are not marked but belong to the genus *Populus*).

2. Role of Culture

Older cross-cultural research on categorization documented apparent deficits in "primitive thinking" along members of traditional societies. For instance, free sorting tasks performed in traditional societies often revealed a preference to sort on the basis of perceptual or functional characteristics, rather than "more advanced" toxonomic classification (e.g., Greenfield, Reich, & Olver, 1966; Price-Williams, 1962; Suchman, 1966). However, more recently, such interpretations have been seriously questioned; performance of members of traditional cultures more likely reflect different interpretations of the task or of the experimenter's expectations. An anecdote reported by Ciborowski (1980) is particularly enlightening: Conducting an experiment among the Kpelle of Liberia, J. A. Glick found that the Kpelle almost always based free-sorts of objects on perceptual or functional similarities, rather than taxonomic kind. He was told that in the traditional Kpelle sense, this was the "clever" way to do the task. In response to this, Glick asked his subjects to perform that task as a "stupid" Kpelle might do it, and the result was perfect taxonomic sorting. Clearly, the fact that the Kpelle were not originally disposed to taxonomic sorting cannot be taken as evidence of primitive thought.

More recent work in cognitive anthropology documents rich systems of knowledge in a variety of domains. For example, ethnobiologists have documented the folk-biological classification systems of members of traditional cultures (e.g., Atran, 1990; Berlin, 1992). Closer examination of conceptual structures in different cultures reveals deep commonalities. Patterns of nomenclature, hierarchical levels of classification, and correspondence of various folk categories to scientific classes show remarkable commonalities across cultures and geographical regions.

One important task remaining for research is to carefully document similarities and differences among conceptual systems of members of different cultures, and among traditional cultures and members of modern, industrialized states. Commonalities may signal deep cognitive universals or a standard interpretation of information in the physical environment (see Coley, Medin, & Atran, 1997). Differences may simply reflect expertise, or different explanatory frameworks, akin to those that Carey (1985) argues characterize adults' and children's conceptions of biology. Either way, careful interdisciplinary research on cultural influences on conceptual systems promises to be interesting, valuable, and extremely difficult.

E. Domain Specificity

Another important trend in recent research on concepts and categorization is the notion of domain specificity. Although there are many ways to characterize the specifics of this proposal, or what counts as a domain (e.g., Hirschfeld & Gelman, 1994; Wellman & Gelman, 1992), what this position boils down to is that all concepts are not necessarily created equal. Understanding concepts necessitates understanding the content of those concepts, and the theoretical framework in which they occur.

1. Biology versus Psychology

One current debate in the study of conceptual development is whether children's understanding of living kinds is embedded in an autonomous biological explanatory theory (Coley, 1995; Inagaki & Hatano, 1993; Keil, 1992, 1994; Springer & Keil, 1989, 1991) or whether concepts of living things are instead imbedded in a protheory that conflates adult notions of biology and psychology (Carey, 1985, 1988, 1995; Solomon, Johnson, Zaitchik, & Carey, 1996). Although the details of the evidence on either side of the issue are beyond the scope of this chapter, the important point is that the meaning of a concept such as *living thing* depends on what framework theory that concept plays a part in.

2. Social Categories

Many of the issues we have been discussing carry over into the social domain (e.g., Smith & Zarote, 1992). With respect to domain specificity, one question is whether people naturalize social categories (Rothbart & Taylor, 1993) or whether social categorization is an autonomous domain with its own constraints and organizing principles (Hirschfeld, 1993, 1994, 1995). Hirschfeld has evidence that the category of race is socially constructed and develops differentially as a function of the ethnic

ɔn of conceptual structures in different cul-
erns of nomenclature, hierarchical levels of
various folk categories to scientific classes
cultures and geographical regions.
esearch is to carefully document similarities
ems of members of different cultures, and
ers of modern, industrialized states. Com-
iversals or a standard interpretation of infor-
e Coley, Medin, & Atran, 1997). Differences
t explanatory frameworks, akin to those that
lts' and children's conceptions of biology.
earch on cultural influences on conceptual
able, and extrêmely difficult.

earch on concepts and categorization is the
h there are many ways to characterize the
its as a domain (e.g., Hirschfeld & Gelman,
this position boils down to is that all con-
Understanding concepts necessitates under-
s, and the theoretical framework in which

nceptual development is whether children's
lded in an autonomous biological explana-
atano, 1993; Keil, 1992, 1994; Springer &
of living things are instead imbedded in a
of biology and psychology (Carey, 1985,
ik, & Carey, 1996). Although the details of
are beyond the scope of this chapter, the
a concept such as *living thing* depends on
ays a part in.

sing carry over into the social domain (e.g.,
domain specificity, one question is whether
ibart & Taylor, 1993) or whether social cat-
ith its own constraints and organizing prin-
lirschfeld has evidence that the category of
s differentially as a function of the ethnic

diversity of children's social environments. For example, consider children's reasoning about the skin color of a child born to mixed race, European-American, African-American parents. Young children give the normatively correct answer of intermediate skin color, but by the sixth grade children living in relatively nonintegrated environments shift to the normatively incorrect view that the child will be have skin color matching that of the African-American parent (Hirschfeld, 1993, 1994). The implications of these sorts of studies remain to be worked out, but their relevance to theories of categorization in general and social cognition in particular can hardly be doubted.

F. Summary

In this section, we have outlined some promising avenues for future research in concepts and categorization. For one, studies of object perception and category representation appear to have much to learn from each other. In general, increasing interaction across subareas would be salutary. Another tend we'd like to see continue is the examination of other conceptual functions beyond categorization; we mentioned conceptual combination and induction to name just two, but many more await more detailed exploration. Studies of the impact of language, culture, and expertise, all promise to enrich and challenge our current notions of conceptual structure and function. And finally, examination of concepts in different domains promises to add even more complexity to the picture. Indeed, none of these avenues of research promises to simplify our portrait of human categorization. However, we feel that they will all help bring it into sharper focus.

VIII. CONCLUSIONS

The field of categories and concepts is very much a work in progress. We anticipate that the twists and turns mentioned at the beginning of this chapter will continue. However, when the field is viewed from a larger time scale, it is easy to see that advances are being made on empirical, methodological, and theoretical fronts. We hope and expect that the present trends will continue: a mutually beneficial interplay between work with natural and with artificial categories, attention to multiple conceptual functions, and greater diversity and interaction with respect to subject populations, kinds of categories, and theoretical perspectives. We hope that reviewers at the end of the next century will be able to refer not to a short history but a *rich* history.

Acknowledgments

This work was partially supported by NIH Grant MH 55079 to Douglas Medin. Sandra Waxman and Edward Wisniewski provided helpful comments on this paper.

References

Ahn, W., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science, 16,* 81–121.

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 50–71.

Atran, S. (1990). *Cognitive foundations of natural history: Towards an anthropology of science.* Cambridge, UK: Cambridge University Press.

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition, 11,* 211–227.

Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure of categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 629–654.

Barsalou, L. W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76–121). Cambridge, UK: Cambridge University Press.

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, M. A. Conway, & P. E. M. Morris (Eds.), *Theories of memories* (pp. 29–101). Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, L. W., Solomon, K., & Wu, L.-L. (in press). Perceptual simulation in conceptual tasks. In M. K. Hiraga, C. Sinha, & S. Wilcox (Eds.), *Cultural, typological and psychological perspectives in cognitive linguistics: The proceedings of the 4th conference of the International Cognitive Linguistics Association, Vol. 3.* Amsterdam: John Benjamins.

Bedard, J., & Chi, M. T. H. (1992). Expertise. *Current Directions in Psychological Science, 1,* 135–139.

Bellezza, F. S. (1984). Reliability of retrieval from semantic memory: Noun meanings. *Bulletin of the Psychonomic Society, 22,* 377–380.

Berlin, B. (1992). *Ethnobiological classification: Principles of categorization of plants and animals in traditional societies.* Princeton, NJ: Princeton University Press.

Biederman, I. (1985). Recognition-by-components: A human image understanding. *Psychological Review, 94,* 115–147.

Billman, D., & Knutson, J. F. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 458–475.

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

Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization* (pp. 141–174). Cambridge, UK: Cambridge University Press.

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

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: Bradford Books.

Carey, S. (1988). Conceptual differences between children and adults. *Mind and Language, 3,* 167–183.

Carey, S. (1995). On the origins of causal understanding. In A. Premack (Ed.), *Causal understanding in cognition and culture* (pp. 268–308). New York: Oxford University Press.

Coley, J. D. (1995). Emerging differentiation of folkbiology and folkpsychology: Attributions of biological and psychological properties to living things. *Child Development, 66,* 1856–1874.

Coley, J. D., Medin, D. L., and Atran, S. (1997). Does rank have its privilege? Inductive inferences within folkbiological taxonomics. *Cognition, 64,* 73–112.

nodel of category construction. *Cognitive Science, 16,*

ecision rules in categorization: Contrasting novice and
nental Psychology: Human Perception and Performance, 18,*

tory: Towards an anthropology of science.* Cambridge, UK:

y & Cognition, 11,* 211–227.
d frequency of instantiation as determinants of graded
ental Psychology: Learning, Memory, and Cognition, 11,*

d its implications for inter-concept similarity. In S. Vos-
alogical reasoning* (pp. 76–121). Cambridge, UK: Cam-

inguistic vagary in concepts: Manifestations of a com-
A. C. Collins, S. E. Gathercole, M. A. Conway, &
pp. 29–101). Hillsdale, NJ: Lawrence Erlbaum Associ-

press). Perceptual simulation in conceptual tasks. In
*Cultural, typological and psychological perspectives in cogni-
ence of the International Cognitive Linguistics Association,*

rrent Directions in Psychological Science, 1,* 135–139.
emantic memory: Noun meanings. *Bulletin of the Psy-*

iples of categorization of plants and animals in traditional*
Press.
ts: A human image understanding. *Psychological Review,*

concept learning and value systematicity: A complex
erimental Psychology: Learning, Memory, and Cognition,*

cepts. *Psychological Review, 77,* 546–556.
tion and memory for instances. In E. Rosch & B. B.
169–211). Hillsdale, NJ: Erlbaum.
egorization: The role of prior processing episodes. In
opment: The ecological and intellectual factors in categoriza-*
idge University Press.
84). Evaluation of exemplar-based generalization and
urnal of Experimental Psychology: Learning, Memory, and*

mbridge, MA: Bradford Books.
children and adults. *Mind and Language, 3,* 167–183.
anding. In A. Premack (Ed.), *Causal understanding in*
: Oxford University Press.
olkbiology and folkpsychology: Attributions of bio-
things. *Child Development, 66,* 1856–1874.
es rank have its privilege? Inductive inferences within
-112.

Chi, M. T. H. (1992). Conceptual change within and across ontological categories: Examples from learning and discovery in science. In R. Giere (Ed.), *Cognitive models of science: Minnesota studies in the philosophy of science* (pp. 129–186). Minneapolis, MN: University of Minnesota Press.

Chi, M. T. H., & Bjork, R. (1991). Modeling expertise. In D. Druckman & R. A. Bjork (Eds.), *In the mind's eye: Enhancing human performance.* Washington, DC: National Academy Press.

Ciborowski, T. (1980). The role of context, skill and transfer in cross-cultural experimentation. In H. C. Triandis & J. W. Berry (Eds.), *Handbook of cross-cultural psychology: Vol. 2. Methodology* (pp. 279–296). Boston: Allyn & Bacon.

Estes, W. K. (1986a). Memory storage and retrieval processes in category learning. *Journal of Experimental Psychology: General, 115,* 155–175.

Estes, W. K. (1986b). Array models for category learning. *Cognitive Psychology, 18,* 500–549.

Estes, W. K. (1994). *Classification and cognition.* Oxford: Oxford University Press.

Farah, M. J., & Kosslyn, S. M. (1982). Concept development. In H. W. Reese & L. Lipsett (Eds.), *Advances in child development and behavior, Vol. 16.* New York: Academic Press.

Flannagan, M. J., Fried, L. S., & Holyoak, J. K. (1986). Distributional expectations and the induction of category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 241–256.

Fried, L. S. & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 234–257.

Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. *Cognitive Psychology, 16,* 341–370.

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition, 23,* 183–209.

Gelman, S. A., & Medin, D. L., (1993). What's so essential about essentialism? A different perspective on the interaction of perception, language, and conceptual knowledge. *Cognitive Development, 8,* 157–167.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings of the nonobvious. *Cognition, 38,* 213–244.

Gelman, S. A., Coley, J. D., & Gottfried, G. M. (1994). Essentialist beliefs in children: The acquisition of concepts and theories. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind.* pp. 341–367. Cambridge, England: Cambridge University Press.

Gentner, D., & Markman, A. B. (1994) Structural alignment in comparison: No difference without similarity. *Psychology Science, 5,* 152–158.

Goldstone, R. L. (1994a). The role of similarity in categorization: Providing a groundwork. *Cognition, 52,* 125–157.

Goldstone, R. L. (1994b). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20,* 3–28.

Goldstone, R. L. (1994c). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123,* 178–200.

Goldstone, R. L. (1995). Mainstream and avant-garde similarity. *Psychological Belgica, 32,* 145–165.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relations attributes and the non-independence of features in similarity judgments. *Cognitive Psychology, 23,* 222–262.

Goldstone, R. L., & Medin, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 29–50.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects* (pp. 437–447). New York: Bobbs-Merrill.

Greenfield, P. M., Reich, L. C., & Olver, R. R. (1966). On culture and equivalence. In J. S. Bruner, R. R. Olver, & P. M. Greenfield (Eds.), *Studies in Cognitive Growth* (pp. 270–318). New York: Wiley.

Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics III: Speech Acts.* New York: Seminar Press.

Hampton, J. A. (1982). A demonstration of intransitivity in natural concepts. *Cognition, 12,* 151–164.

Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition, 15,* 55–71.

Hampton, J. A. (1988). Over extension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 14*, 12–32.
Harvard, S. (Ed.) (1987). *Categorical perception*. Cambridge, England: Cambridge University Press.
Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. *Bulletin of the Psychonomic Society, 28*, 242–247.
Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 411–422.
Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93*, 411–428.
Hintzman, D. L., & Ludlum, G. (1980). Differential forgetting of prototypes and old instances: Simulation by an exemplar-based classification model. *Memory & Cognition, 8*, 378–382.
Hirschfeld, L. A. (1993). Discovering social difference: The role of appearance in the development of racial awareness. *Cognitive Psychology, 25*, 317–350.
Hirschfeld, L. A. (1994). Is the acquisition of social categories based on domain-specific competence or on knowledge transfer? In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 201–233). Cambridge, England: Cambridge University Press.
Hirschfeld, L. A. (1995). Anthropology, psychology and the meanings of social causality. In A. Premack (Ed.), *Causal understandings in cognition and culture* (pp. 313–350). New York: Oxford University Press.
Hirschfeld, L. A., & Gelman, S. A. (Eds.) (1994). *Mapping the Mind: Domain specificity in cognition and culture.* New York: Cambridge University Press.
Hock, H. S., Webb, E., & Cavedo, L. C. (1987). Perceptual learning in visual category acquisition. *Memory & Cognition, 15*, 544–556.
Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295–355.
Homa, D. (1984). On the nature of categories. In G. Bower (Ed.), *The Psychology of Learning and Motivation, 18*, 49–94.
Homa, D., & Chambliss, D. (1975). The relative contributions of common and distinctive information on the abstraction from ill-defined categories. *Journal of Experimental Psychology: Human Learning and Memory, 1*, 351–359.
Homa, D., Cross, J., Cornell, D., Goldman, D., & Schwartz, S. (1973). Prototype abstraction and classification of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology, 101*, 116–122.
Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 418–439.
Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99*, 480–517.
Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind-body distinction. *Child Development, 64*, 1534–1549.
Jones, S. S., & Smith, L. B. (1993). The place of perception in children's concepts. *Cognitive Development, 8*, 113–139.
Katz, J. J., & Postal, P. M. (1964). *An integrated theory of linguistic descriptions.* Cambridge, MA: MIT Press.
Keil, F. C. (1979). *Semantic and conceptual development: An ontological perspective.* Cambridge, MA: Harvard University Press.
Keil, F. C. (1981). Constraints on knowledge and cognitive development. *Psychological Review, 88*, 197–227.
Keil, F. C. (1986). The acquisition of natural kind and artifact terms. In W. Demopoulos & A. Marras (Eds.), *Language Learning and Concept Acquisition* (pp. 133–153).
Keil, F. C. (1989). *Concepts, kinds, and cognitive development.* Cambridge, MA: MIT Press.

Keil, F. C. (1992). The origins of an autonomous biology. In M. A. Gunnar & M. Maratsos (Eds.) *Minnesota Symposium on Child Psychology, Vol. 25,* (pp. 103–138). Hillsdale, NJ: Erlbaum.

Keil, F. C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind* (pp. 234–254). Cambridge: Cambridge University Press.

Kemler-Nelson, D. G. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior, 23,* 734–759.

Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin, 112,* 500–526.

Kruschke, J. K. (1993). Three principles for models of category learning. In G. V. Nakamura, R. Taraban, & D. L. Medin (Eds.), *The Psychology of Learning and Motivation: Categorization by humans and machines,* vol. 29. pp. 57–90. San Diego: Academic Press, Inc.

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

Kurbat, M. (1995). Part-based models of visual similarity, typicality, and categorization. Doctoral dissertation, Department of Psychology, University of Michigan.

Kurbat, M., Smith, E. E., & Medin, D. L. (1994). Shape similarity, typicality and categorization. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago: University of Chicago Press.

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

Lassaline, M. E., Wisniewski, E. J., & Medin, D. L. (1992). Basic levels in artificial and natural categories: Are all basic levels created equal? In B. Burns (Ed.), *Percepts, concepts, and categories: The representation and processing of information* (pp. 323–378). New York: Elsevier.

Levine, M. (1975). *A cognitive theory of learning: Research on hypothesis testing.* Hillsdale, NJ: Erlbaum.

Malt, B. C. (1990). Features and beliefs in the mental representations of categories. *Journal of Memory and Language, 29,* 289–315.

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

Malt, B. C., & Smith, E. E. (1983). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23,* 250–269.

Markman, A. B., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language, 32,* 517–535.

Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431–467.

Markman, A. B., Yamauchi, T., & Makin, V. S. (1997). The creation of new concepts: A multifaceted approach to category learning. To appear in T. B. Ward, S. M. Smith, & J. Vaid (Eds.), *Conceptual Structures: Emergence, Discovery, and Change* (pp. 179–208). Washington, DC: American Psychological Association.

Markman, E. M. (1987). How children constrain the possible meanings of words. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization,* pp. 256–287.

Martin, R. C. & Caramazza, A. (1980). Classification of well-defined and ill-defined categories: Evidence for common processing strategies. *Journal of Experimental Psychology: General, 109,* 320–353.

Massey, C. M., & Gelman, R. (1988). Preschoolers' ability to decide whether a photographed unfamiliar object can move itself. *Developmental Psychology, 24,* 307–317.

Matheus, C. J., Rendell, L. R., Medin, D. L., & Goldstone, R. C. (1989). Purpose and conceptual functions: A framework for concept representation and learning in humans and machines. In *Proceedings of the Seventh Annual Conference of the Society for the Study of Artificial Intelligence and Simulation of Behavior.*

Mayr, E. (1988). *Toward a New Philosophy of Biology.* Cambridge, MA: Belknap Press of Harvard University Press.

McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? *Memory and Cognition, 6,* 462–472.

McNamara, T. P., & Sternberg, R. J. (1983). Mental models of word meaning. *Journal of Verbal Learning and Verbal Behavior, 22,* 449–474.

Medin, D. L. (1986). Comment on "Memory storage and retrieval processes in category learning." *Journal of Experimental Psychology: General, 115,* 373–381.

Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist, 44,* 1469–1481.

Medin, D. L., & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience, *Journal of Experimental Psychology: General, 117,* 68–85.

Medin, D. L., & Heit, E. J. (in press). Categorization. In D. Rumelhart & B. Martin (Eds.), *Handbook of cognition and perception.* Hillsdale, NJ: Erlbaum.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning,* 189–223.

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving and induction. In R. J. Sterberg (Ed.), *Advances in the Psychology of Human Intelligence, 5,* 189–223.

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

Medin, D. L. & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 355–368.

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

Medin, D. L., & Smith, E. E. (1981). Strategies and classification learning. *Journal of Experimental Psychology: Human learning and Memory, 7,* 241–253.

Medin, D. L., & Wattenmaker, W. D. (1987). Category cohesiveness, theories, and cognitive archeology. In U. Neisser (Ed.), *Concepts and Conceptual Development: The Ecological and Intellectual Factors in Categories* (pp. 25–62). Cambridge, England: Cambridge University Press.

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

Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 607–625.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100,* 254–278.

Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processes and decision processes. *Psychonomic Bulletin & Review, 2,* 1–19.

Medin, D. L., Lynch, E. B., Coley, J. D., Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology, 32,* 49–96.

Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, concept cohesiveness, and category construction. *Cognitive Psychology, 19,* 242–279.

Medin, D. L., Wattenmaker, W. D., & Michalski, R. S. (1987). Constraints and preferences in inductive learning: An experimental study of human and machine performance. *Cognitive Science, 11,* 299–339.

Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual Review of Psychology, 32,* 89–115.

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

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

Murphy, G. L., & Wisniewski, E. J. (1989). Categorizing objects in isolation and in sciences: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 572–586.

Nakamura, G. V., Taraban, R., & Medin, D. L. (Eds.). (1993). *The Psychology of Learning and Motivation: Categorization by humans and machines,* vol. 29. San Diego: Academic Press, Inc.

ıtal models of word meaning. *Journal of Verbal Learning*

rage and retrieval processes in category learning." *Jour-*
373–381.
ructure. *American Psychologist, 44,* 1469–1481.
ı structure and the use of base-rate information from
y: General, 117, 68–85.
ion. In D. Rumelhart & B. Martin (Eds.), *Handbook of*
m.
essentialism. In S. Vosniadou & A. Ortony (Eds.), *Sim-*

: character of abstract thought: Categorization, prob-
ʒ (Ed.), *Advances in the Psychology of Human Intelligence,*

theory of classification learning. *Psychological Review,*

r separability in classification learning. *Journal of Exper-*
·ory, 7, 355–368.
d structure in conceptual combination. *Cognitive Psy-*

nd classification learning. *Journal of Experimental Psy-*
253.
ʒory cohesiveness, theories, and cognitive archeology.
Development: The Ecological and Intellectual Factors in Cat-
ımbridge University Press.
reko, D. (1982). Correlated symptoms and simulated
Psychology: Learning, Memory, and Cognition, 8, 37–50.
·). Relationships between item and category learning:
urnal of Experimental Psychology: Learning, Memory, and

'93). Respects for similarity. *Psychological Review, 100,*

. (1995). Comparison and choice: Relations between
sychonomic Bulletin & Review, 2, 1–19.
ꜱ. (1997). Categorization and reasoning among tree
e Psychology, 32, 49–96.
E. (1987). Family resemblance, concept cohesiveness,
gy, 19, 242–279.
.. S. (1987). Constraints and preferences in inductive
ı and machine performance. *Cognitive Science, 11,*

ı of natural objects. In M. R. Rosenzweig & L. W.
, 89–115.
oncepts. *Cognitive Science, 12,* 529–562.
heories in conceptual coherence. *Psychological Review,*

orizing objects in isolation and in sciences: What a
ntal Psychology: Learning, Memory, and Cognition, 15,

ꜱ.). (1993). *The Psychology of Learning and Motivation:*
San Diego: Academic Press, Inc.

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 87–108.

Nosofsky, R. M. (1988a). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 700–708.

Nosofsky, R. M. (1988b). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 54–65.

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition in memory. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 3–27.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Erlbaum.

Nosofsky, R. M. & Kruschke, J. K. (1992). Investigations of an exemplar-based connectionist model of category learning. In D. L. Medin (Ed.) *The Psychology of Learning and Motivation* (Vol. 28) (pp. 207–250). New York: Academic Press.

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

Nosofsky, R. M., Bluck, M. A., Palmeri, T. S., McKinley, S. C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition, 22,* 352–369.

Nosofsky, R. M., Palmeri, T. J., & McKinely, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101,* 53–79.

Oden, G. C. (1987). Concept, knowledge, and thought. *Annual Review of Psychology, 38,* 203–227.

Ortony, A., Vondruska, R. J., Foss, M. A., & Jones, L. E. (1985). Salience, similes, and the asymmetry of similarity. *Journal of Memory and Language, 24,* 569–594.

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

Peterson, M. J., Meagher, R. B., Jr., Chait, H., & Gillie, S. (1973). The abstraction and generalization of dot patterns. *Cognitive Psychology, 4,* 378–398.

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

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

Price-Williams, D. R. (1962). Abstract and concrete models of classification in a primitive society. *British Journal of Educational Psychology, 32,* 50–61.

Reagher, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General, 122,* 92–114.

Reagher, G., & Brooks, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 347–363.

Rips, L. (1995). The current status of research on conceptual combination. *Mind and Language. 10,* 72–104.

Rips, L. J. (1975). Inductive judgements about natural categories. *Journal of Verbal Learning and Verbal Behavior, 14,* 665–681.

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

Rips, L. J., & Collins, A. (1993). Categories and resemblance. *Journal of Experimental Psychology: General, 122,* 468–486.

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

Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439.

Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15*,346–378.

Rothbart, M., & Taylor, M. (1992). Category labels and social reality: Do we view social categories as natural kinds? In G. Semin and K. Fiedler (Eds.), *Language, Interaction, and Social Cognition* (pp. 11–36). London: Sage Publications.

Schank, R. C., Collins, G. C., & Hunter, L. E. (1986). Transcending inductive category formation in learning. *Behavioral and Brain Sciences, 9*, 639–651.

Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. L. Medin (Ed.) *The Psychology of Learning and Motivation, 31*, 305–354. San Diego, CA: Academic Press.

Sebestyn, G. S. (1962). *Decision-making processes in pattern recognition.* New York: Macmillan.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75.*

Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121*, 278–304.

Shipley, E. F. (1993). Categories, hierarchies, and induction. In D. L. Medin (Ed.). *The Psychology of Learning and Motivation, Vol. 30.* (pp. 265–301). New York: Academic Press.

Smith E. E., & Medin, D. L. (1981). Categories and concepts. Cambridge, MA: Harvard University Press.

Smith, E. E., & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science, 8*, 337–361.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane M. (1988). Combining prototypes: A selective modification model. *Cognitive Science, 12*, 485–527.

Smith, E. E., Shafir, E., & Osherson, D. N. (1993). Similarity, plausibility, and judgments of probability. *Cognition, 49*, 67–96.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81*, 214–241.

Smith, E. R., & Zarate, M. A. (1992). Exemplar-based models of social judgment. *Psychological Review, 99*, 3–21.

Solomon, G. E. A., Johnson, S. C., Zaitchik, D., & Carey, S. (1996). Like father, like son: Young children's understanding of how and why offspring resemble their parents. *Child Development, 67*, 151–171.

Spalding, T. L. & Murphy, G. L. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory & Cognition, 22*, 525–538.

Springer, K., & Keil, F. C. (1989). On the development of biologically specific beliefs: The case of inheritance. *Child Development, 60*, 637–648.

Springer, K., & Keil, F. C. (1991). Early differentiation of causal mecanisms appropriate to biological and nonbiological kinds. *Child Development, 62*, 767–781.

Suchman, R. G. (1966). Cultural differences in children's color and form perception. *Journal of Social Psychology, 70*, 3–10.

Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology, 23*, 457–482.

Tarr, M. J. (1995). Rotating objects to recognize them: A case study of the role of mental transformations in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review, 2*, 55–82.

Thibaut, J.-P., & Schyns, P. G. (1995). The development of feature spaces for similarity and categorization. *Psychological Belgica, 35*, 167–185.

Trabasso, T., & Bower, G. H. (1968). *Attention in learning.* New York: Wiley.

Tulving, E. (1983). *Elements of episodic memory.* New York: Oxford University Press.

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

VanMechlen, I., Hampton, J., Michalski, R., & Theuns, P. (Eds.). (1993). *Categories and Concepts: Theoretical Views and Inductive Data Analysis.* London: Academic Press.

Watanabe, S. (1969). *Knowing and guessing: A formal and quantitative study.* New York: Wiley.

of context on the structure of categories. *Cognitive Psy-*

ıbels and social reality: Do we view social categories as
·dler (Eds.), *Language, Interaction, and Social Cognition*

(1986). Transcending inductive category formation in
39–651.

geny of part representation in object concepts. In D. L.
d Motivation, 31, 305–354. San Diego, CA: Academic

n pattern recognition. New York: Macmillan.
. (1961). Learning and memorization of classifications.

-scaling studies of dot-pattern classification and recog-
ieneral, 121, 278–304.

nduction. In D. L. Medin (Ed.). *The Psychology of Learn-*
New York: Academic Press.

.d concepts. Cambridge, MA: Harvard University Press.

ıal combination with prototype concepts. *Cognitive Sci-*

ıe M. (1988). Combining prototypes: A selective mod-
527.

). Similarity, plausibility, and judgments of probability.

Structure and process in semantic memory: A featural
Review, 81, 214–241.

based models of social judgment. *Psychological Review,*

& Carey, S. (1996). Like father, like son: Young chil-
pring resemble their parents. *Child Development, 67,*

ɔf background knowledge on category construction.
Memory & Cognition, 22, 525–538.

nent of biologically specific beliefs: The case of inher-

:ion of causal mecanisms appropriate to biological and
767–781.

ldren's color and form perception. *Journal of Social Psy-*

ries and expertise: Is the basic level in the eye of the

hem: A case study of the role of mental transforma-
ıl objects. *Psychonomic Bulletin & Review, 2,* 55–82.

ɔment of feature spaces for similarity and categoriza-

arning. New York: Wiley.

v York: Oxford University Press.

ical Review, 84, 327–352.

.euns, P. (Eds.). (1993). *Categories and Concepts: Theo-*
ɔn: Academic Press.

and quantitative study. New York: Wiley.

Wattenmaker, W. D. (1991). Learning modes, feature correlations, and memory-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 908–923.

Wattenmaker, W. D. (1993). Incidental concept learning, feature frequency and correlated properties. *Journal of Experimental Psychology: Human learning and memory, 19,* 203–222.

Wattenmaker, W. D., Dewey, G. I., Murphy, T. D., Medin, D. L. (1986). Linear separability and concept learning: Context, relational properties, and concept naturalness. *Cognitive Psychology, 18,* 158–194.

Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development, 5*(2), 123–150.

Waxman, S. R., & Gelman, R. (1986). Preschoolers' use of superordinate relations in classification and language. *Cognitive Development, 1,* 139–156.

Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence form 12-month-old infants. *Cognitive Psychology, 29,* 257–302.

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

Wisniewski, E. J. (1995). Prior Knowledge and functionally relevant features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 449–468.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory & Language, 35,* 1–20.

Wisniewski, E. J. (1997). When concepts combine, *Psychonomic Bulletin & Review, 4,* 167–183.

Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 241–284). Amsterdam: North Holland.

Wisniewski, E. J., & Medin, D. L. (1991). Harpoons and long sticks: The interaction of theory and similarity in rule induction. In D. H. Fisher, M. J. Pazzani, & P. Langley (Eds.), *Concept formation: Knowledge and experience in unsupervised learning.* pp. 237–278. San Mateo, CA: Morgan Kaufman.

Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science, 18,* 221–281.