2.29 Concept and Category Learning in Humans

B. H. Ross, E. G. Taylor, and E. L. Middleton, University of Illinois at Urbana-Champaign, Urbana-Champaign, IL, USA

T. J. Nokes, University of Pittsburgh, Pittsburgh, PA, USA

2.29.1 Introduction

Concepts and categories are crucial for intelligent thought and action. A child needs to learn to tell toys from tools and which types of dogs can be petted. A student needs to learn to distinguish the principle underlying a math problem, so that relevant principle knowledge can be applied. Researchers need to be able to decide what type of person has asked them to collaborate, to understand the concept of confounding, and to be able to communicate new ideas to others. The focus of this chapter is understanding the learning of categories and concepts. This learning is critical not only because little of this knowledge is innate but also because the learning of concepts and categories is a large part of acquiring knowledge.

We use concept to refer to a mental representation, an idea, that picks out a class of entities. A category is the class or set of entities that is referred to by the concept. One has a concept of dog, and then there is the category of dogs (both real and imagined). Different things in the category can all be treated similarly, with respect to some purpose. Cars are members of a category because they have much in common in terms of appearance and use. However, one might want to make finer distinctions, such as separating cars that can carry much stuff for moving from sports cars; at other times, such as in compiling assets, one might want to think of an expensive sports car as part of the same class as a boat and house. When possible, we use categories to refer to the learning or use of the members and concepts to refer to the knowledge of categories. Despite this distinction, they are often used interchangeably because it is often true that one is dealing with both together.

We also need to address briefly two important issues: Origins and types. First, although some researchers argue for the independent existence of concepts in the world (separate from the organisms that perceive them), we, as psychologists, believe it is more useful to think of them as arising jointly from the fit between the world and the organism. Although it is true that we do not just have any concepts, it is also true that different organisms have different concepts, and even among humans, the concepts we have often are a function of human activity. Second, although much of the work has focused on object categories, we clearly have many other categories besides objects, people, situations, problems, scenes, and so on. We try to include a variety of types in our discussion (also see Medin et al., 2000).

In this chapter, we provide an overview of the work on concept and category learning, with a focus on experimental work and modeling. We begin by considering the functions of concepts. We then address a large body of research that has examined conceptual structure for classification, with particular emphasis on prototype and exemplar approaches and models. We argue that this work has missed some important aspects of both how categories are learned and the importance of structure and prior knowledge. A goal of this chapter is to integrate the work on concepts and categories into the areas of cognition in which they are so crucial. All cognitive activity relies on concepts in some way, so an understanding of how they are learned is likely to have an impact on much other research. We provide an integration in three ways. One, we examine the learning of concepts and categories from a goal-oriented view, asking how the ways in which they might be learned influence the representation available for other cognitive activities. Two, we consider more complex concepts, as would be needed to account for results in most areas in which concepts are used. Three, we address this integration more directly by considering two very different areas of cognition, problem solving and language.

2.29.1.1 Functions of Concepts and Categories

Concepts and categories are fundamental building blocks of cognition. Murphy (2002) calls them the mental glue in that they link our past experiences (with toys, dogs, mathematical problems, collaborators) to our current ones. They are a part of all cognitive, and many noncognitive, theories. We mention here a few functions.

2.29.1.1.1 *Classification*

Classification is the determination that something is a member of a particular category: A carrot, an extrovert, a permutations problem, an instance of insurgency. This action allows one to access knowledge about the category that can be used for a variety of other functions.

2.29.1.1.2 *Prediction and inference*

When an object is classified as a carrot, you can predict how it will taste, how crunchy it will be. You can use knowledge of the category to infer how it was grown and how similar it is to a beet. Prediction is often considered a key function of categories (e.g., Anderson, 1991) in that it allows for selection of plans and actions.

2.29.1.1.3 *Understanding and explanation*

People need to know not just what, but why. We can explain aberrant behavior if we know that the person was grieving or drunk. This characterization might change our future actions with the person. If we start watching a TV program part-way through and cannot understand what is happening, being told it involves a love triangle may help us to make sense of the events.

2.29.1.1.4 Communication

We are social animals, and much of our activity is geared toward interacting with others. Concepts provide a kind of social glue as well, as they facilitate communication and allow us to learn new concepts by indirect experience.

Concepts and categories underlie much of mental life, and their learning is complex. We must learn not just to classify items: Knowing something is a carrot, a permutations problem, or an extrovert does not help unless we know enough about carrots, permutations, or extroverts to accomplish our goals (e.g., how to prepare the carrot, solve the problem, quiet the extrovert). Thus, as we learn to tell what category an item is in, we also have to learn much else about the concept to allow prediction, explanation, and communication. The corresponding distinction between knowledge that allows one to classify and knowledge that allows one to perform other conceptual functions is a central one for this chapter, as well as for understanding concepts and categories in other areas of cognition.

2.29.2 Conceptual Structure

We turn now to the structure of concepts. How are categories mentally represented to allow these various functions to be accomplished? This question is essential for any examination of concept and category learning, because it provides a clear target for what must be learned. We can only provide a brief overview, but fuller descriptions can be read in Medin (1989) and Murphy (2002).

2.29.2.1 Views and Models

What determines which items go together in a category and which items are in different categories? A common intuition is that it depends upon similarity – more similar items are in the same category. One can think of items as consisting of a set of features. Similarity is defined as the overlap of features (such as Tversky, 1977) or, if one prefers spatial metaphors, as closeness in some multidimensional space (e.g., Shepard, 1962a,b). This idea of similarity underlies many views of conceptual structure, three of which are summarized in **Table 1**. We briefly present these views and associated formal models with a focus on classification learning.

2.29.2.1.1 Classical view

The classical view of concepts takes a strict view of similarity: All items in a category must have a specific set of features. If an item has those features, it is in the category; if it does not, it is not. One can think of this as a definitional view of categories: The features are singly necessary and jointly sufficient for category membership. A triangle is any closed two-dimensional figure that has three straight sides. Any item that has all of those features is a triangle, and any item that does not have all of those features is not a triangle. In addition, the view includes the rule-based idea, in which items are classified as being in a category if they meet some rule, such as red, or red and large (e.g., Bruner et al., 1956). This view has a long history (see Murphy, 2002), and it matches many intuitions about category members sharing some common characteristics. However, because the classical view assumes all members possess the same set of common features, it does not explain why some category members are more typical than others (e.g., robins vs. penguins) or why it has proven so difficult for people to come up with a set of defining features for most categories (Wittgenstein, 1953, has a famous example of trying to define the category games). There are no current formal models that rely solely on a classical view.

2.29.2.1.2 Prototype view

The prototype view (or probabilistic view) keeps the attractive assumption that there is some underlying common set of features for category members but relaxes the requirement that every member have all the features. Instead, it assumes there is a probabilistic matching process: Members of the category have more features, perhaps weighted by importance, in common with the prototype of this category than with

Table 1 Similarity-based views of conceptual structure and their classification decisions

Classical	Unitary description: Definitional, rule-based	Classification: Category member if and only if all features are true of an item
Prototype	Unitary description (prototype): Probabilistic	Classification: Category member if more (weighted) features are true of the prototype than of other prototypes
Exemplar	No unitary description: Disjunctive representation	Classification: Category member if more similar to (weighted) category members than to members of other categories

prototypes of other categories (and perhaps some minimum match level is required). An early presentation is available in Rosch and Mervis (1975), with a more recent presentation in Hampton (2006).

This type of representation has major implications for how to think about categories. First, some members may have more of the prototype features than others, such as a robin having more bird prototype features than a penguin. People tend to judge robins as better examples of birds than are penguins and are faster to classify robins as birds than they are to classify penguins as birds. Second, this view suggests that the category boundaries, which were strict in the classical view, may be fuzzy, with some cases that are far from the prototype and maybe even almost as close to another prototype. For example, whales have many fish-like properties, and bats have many bird-like properties. Both are viewed as poor examples of the mammal category and are slow to be verified as members of that category. Overall, this view leads to a set of category members that tend to have a family resemblance – no defining features, but some features will be possessed by many members and some features by a few members (similar to an extended family). Rosch and Mervis (1975) argue that this type of family resemblance characterizes many natural categories. See the top half of **Figure 1** for a simple example.

Smith and Minda (1998, 2000; Minda and Smith, 2001) proposed a model of prototype-based classification that matches an item to the various prototype representations and picks the most similar (see also Hampton, 1993). Using the spatial distance idea of similarity, the prototype models (a) determine the distance from the test item to each prototype, (b) compute the similarity from the distance, and (c) choose the category prototype as a function of the similarity. There are several specific choices to make as to how to formalize these ideas. For the reader interested in formal models, we provide a simple one with additive similarity and just two categories in **Table 2**.

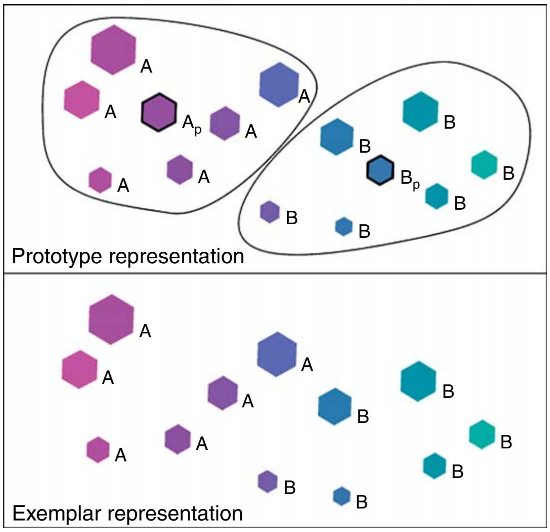

Figure 1 Sample prototype and exemplar representations. Members of Category A tend to be large and pink, whereas members of Category B tend to be small and turquoise. The prototype representation in the top panel also includes a prototype item with the bold outline and P subscript, which was not presented but represents the center of the presented items. (Although the prototypes here can be characterized as instances, more generally prototypes are summary representations.) The exemplar representations in the bottom panel include all individual category members, but do not group them or include a center.

The learning of the category is the building of this summary representation. Exactly how the prototype is learned is not usually specified. One possibility, for simple cases, is to assume a simple associationist-like mechanism, with frequently occurring values being more reinforced. The point is that the earlier experiences are used to build the summary representation,

Table 2 Sample prototype model calculations (assuming *N* dimensions, city-block[a] metric, and two categories, A and B)

To calculate	*Formula*	*Comment*
(a) Distance of test item, *t*, to prototype, *P*	$d_{tP} = \sum_{k=1}^{N} w_k \lvert t_k - P_k \rvert$	w_k Weights each dimension *k* by attention (or importance)[b]
(b) Similarity of test item, *t*, to prototype, *P*	$S_{tP} = 1 - d_{tP}$	Additive similarity [c]
(c) Probability of choosing Category *A*	$P(A \mid t) = \dfrac{S_{tP,A}}{S_{tP,A} + S_{tP,B}}$	Luce Choice rule over prototypes' similarities to the test item

[a]City-block metric is common when the dimensions' distances seem to be evaluated separately and then added. It is as if one is walking in a city and can turn only at the corners, rather than as the crow flies (Euclidean). It is used here as an illustration of one distance metric.
[b]The difference between the values of the test item and prototype on each dimension is summed, weighted by the attention given to that dimension. The sum of the weights is constrained to be equal to 1 (Σ $w_k = 1$) as a limited-capacity system (total amount of attention is limited).
[c]Similarity is assumed to be a linearly decreasing function of psychological distance.

which is then slightly revised with new experiences (in the same way one can keep various statistics of a sample as each new item is added).

2.29.2.1.3 Exemplar view

The idea of a summary representation is attractive, and we clearly know a lot about categories such as birds. However, an alternative view is that the knowledge we use in making category judgments such as classification is not this summary representation but the set of category members: The conceptual structure is a collection of the mental representations of category members. When an item is presented, the person matches the representation of the item to the various previous items' representations and uses some subset (perhaps even just one item) selected by their similarity to the current item's representation. Thus, to classify a hawk as a bird, one would retrieve some earlier representations of hawks or similar birds and note their bird category; to decide what it ate, one would again use these more specific representations.

We know this lack of a summary representation seems strange at first thought, but often a new item will remind one of some old item from which one can make a classification or prediction ("is this person like my friend Bill in other ways?"). The exemplar view is an extension of this idea to allow the use of specific instances to classify the current item. Imagine you see a new ostrich-like animal with some unusual bird characteristics. If you relied on the bird prototype, you might not classify it as a bird. However, if you used similar birds, then you would likely classify it as a bird, as well as predict it probably does not fly or eat worms.

This exemplar view accounts for the prototype-like effects mentioned earlier. Some category members, such as robins, will be viewed as more typical than others, such as penguins, because a robin will be similar to many birds (not just the large number of robins, but also sparrows, thrushes, etc.), whereas a penguin will be similar mainly to the few penguins one has seen. The uncertain cases will be exemplars that are not very similar to any items in a category and may even be a little bit similar to atypical items in another category (see **Figure 1**, lower panel).

There are a number of formal exemplar models, with the best known being the Context Model (Medin and Schaffer, 1978), which was extended in a number of ways by the Generalized Context Model (GCM, Nosofsky, 1986). Like the prototype model, exemplar models also use the similarity of the test item to each category, but because there is no summary representation, the similarity of the test item to the category is just the sum of the similarity of the test item to each of the items in the category. The other main difference from prototype models is that multiplicative, rather than additive, similarity is assumed. Multiplicative similarity means that the similarity is not a linear function of the psychological distance, but instead, very close items (ones that are similar to the test item on many dimensions) matter much more. More specifically, exemplar models (a′) compare the similarity of the test item to all items, (b′) compute the similarity of the test item to each category by summing the test item's similarity to each of the category members, and then (c′) choose the category as a function of similarity. The equations for one such model are presented in **Table 3** for those readers with an interest in formal models. Most exemplar models capture only classification performance, but ALCOVE (Kruschke, 1992) extends the GCM to allow some

Table 3 Sample exemplar model calculations (assuming N dimensions, city-block[a] metric, and two categories, A and B)

To calculate	*Formula*	*Comment*
(a′) Distance of test item, t, to exemplar, X_j	$d_{tj} = \sum_{k=1}^{N} w_k \lvert t_k - X_{jk} \rvert$	w_k Weights each dimension k by attention (or importance)[b]
(b′) Similarity of test item, t, to exemplar, X_j	$S_{tj} = \exp(-c \cdot d_{tj})$	Multiplicative similarity,[c] c indicates sharpness of generalizations
(c′) Probability of choosing Category A	$P(A \mid t) = \dfrac{\sum_{i \in A} S_{ti}}{\sum_{i \in A} S_{ti} + \sum_{i \in B} S_{ti}}$	Luce Choice rule over exemplars' similarities to the test item

[a]City-block metric is common when the dimensions' distances seem to be evaluated separately and then added. It is as if one is walking in a city and can turn only at the corners, rather than as the crow flies (Euclidean). It is used here just as an illustration of one distance metric.
[b]The difference between the values of the test item and prototype on each dimension is summed, weighted by the attention given to that dimension. The sum of the weights is constrained to be equal to 1 ($\Sigma\, w_k = 1$) as a limited-capacity system (total amount of attention is limited).
[c]Multiplicative similarity is not a linear function of psychological distance. Rather, close items (small distance) have much greater effect than in a linear function. The exponential function is commonly used and has a quick drop-off with distance. (Note that the exponential of the sum of distances across each dimension is equal to multiplying the exponential of each dimensional distance, since $\exp[a + b] = \exp[a] \cdot \exp[b]$.)

learning. In addition, there is no learning of summary representations, because only individual items need to be stored.

2.29.2.2 Evaluations of Prototype and Exemplar Models

There have been many (many) comparisons of prototype and exemplar models (for a review, see Murphy, 2002). The general result is that exemplar models do as well or better than prototype models in most cases (though see Smith and Minda, 1998, 2001). Our view, summarizing over many results, is that the advantage is largely a result of two factors: similarity calculation and selectivity (Ross and Makin, 1999). First, the exemplar models' multiplicative similarity (compared to the prototype models' usual assumption of additive similarity) means that the model does not combine features independently but, rather, is sensitive to the relational information among the features (e.g., Medin and Schaffer, 1978). The combination of features is being used beyond their separate contribution to determining classification. Thus, if one encountered small birds that sang and large birds that squawked, a prototype representation would not be sensitive to that particular relational (cooccurrence) information, whereas an exemplar model would. Although it is not a usual assumption, prototype models might also incorporate multiplicative similarity (e.g., Smith and Minda, 1998); this helps the fit, but it does not mimic the predictions of the exemplar model. (Multiplicative similarity is a nonlinear function, so calculating multiplicative similarity on the mean (prototype) is not the same as the mean of multiplicative similarities on the individual instances.) In fact, exemplar models implicitly keep all the statistical information (e.g., frequency, variability, cooccurrence) by keeping all the exemplars. One might argue that a prototype model could also keep various statistical information around to make it equivalent to such exemplar models (e.g., Barsalou, 1990), but no one has proposed such a formal model. Second, because exemplar models have no summary representation, the same knowledge is not used for all the different decisions. Thus, even unusual items can be classified by similarity to earlier unusual similar members. The ability to use different knowledge for different decisions means that the exemplar model can classify unusual items without compromising its ability to easily classify more typical items. This flexibility is important in allowing the exemplar model to fit a variety of classification data.

The exemplar model fits the data well for many classification experiments but has difficulties with other aspects of category-related judgments. One major issue is that it has no place for these summary representations that we all find attractive in thinking about concepts. In particular, to answer the question as to why these items are all members of the same category, the exemplar view is left with the unsatisfying answer that "they all have the same category label." That may be fine for arbitrary experimenter-defined categories in the laboratory but seems woefully inadequate for permutation problems, extroverts, and love triangles. (Note that the classical view can point to a definition and the prototype view to similarity to some common summary representation.)

2.29.2.3 Combined Models

It seems likely that people are not restricted to a single means of representation and that we might combine advantages of prototype and exemplar models, or at least general and more specific knowledge. We consider two types of proposals. First, one could take a prototype model and an exemplar model and simply combine them with some means of determining whether a decision would be based on the prototypes or exemplars. Smith and Minda (1998) show that a combined model can provide a better account of the data, with the prototype being more influential earlier, when each item to be classified has been presented only a few times, and the exemplar model controlling responses more as the same items are presented often. They do not specify a control mechanism, but perhaps the model that has greater confidence in its choice might determine the classification.

Second, there have been models that build upon simple models with a more integrated approach. Interestingly, rules, which can be viewed as simple classical models, are making a comeback: They appear to work better as part of the answer rather than the sole answer. ATRIUM (Erickson and Kruschke, 1998) combines simple rules with an exemplar model, ALCOVE (Kruschke, 1992). The authors provide experimental evidence showing that both types of knowledge can influence a task and then address how the two types of knowledge might be integrated to provide an account of the data. All inputs are examined by both the rule and exemplar modules, with the response a weighted function determined by how much attention is given to each. The model learns to shift attention between the modules as a function of which module is better at classifying particular inputs. (Also see the RULEX model by Nosofsky et al., 1994.)

A very different combined model, COVIS, has been proposed by Ashby and colleagues (Ashby and Ennis, 2006, present an overview). Human category learning is assumed to be mediated by a number of functionally distinct neurobiological systems, and the goal is to elucidate these systems and their behavioral consequences. An explicit system is important for rule-based tasks – those tasks with a focus on a single dimension for which people might generate and test hypotheses. Another, procedural-based, system deals with information-integration tasks that require combining information from multiple dimensions. This model combines rule- and procedural-based knowledge to account for a variety of behavioral results and data from neuropsychological patients.

2.29.2.4 Evaluation of Work on Conceptual Structure

2.29.2.4.1 Successes

The separation of different views of conceptual structure has generated much research. The prototype view has greatly extended our understanding of natural categories. The formal modeling on the exemplar approach has shown that exemplar representations coupled with multiplicative similarity are able to account for a wide variety of classification results. The more recent prototype modeling work shows that some findings that seemed problematic for prototype models may not be, though the exemplar model still seems to have an edge on overall accounts of the results. Although almost all of the exemplar work has focused on learning artificial categories, some recent work suggests the exemplar models may fit some real-world categories as well (reviewed in Storms, 2004).

2.29.2.4.2 Limitations

We label this part of the evaluation 'limitations' because the difficulties are not failures but restrictions. The simple point, to be elaborated in the next section, is that the field has examined only a small part of the picture for conceptual structure and learning. Thus, although the prototype and exemplar approaches and models have been explored extensively, especially in the laboratory, our understanding of concept and category learning may be quite limited.

First, almost all the work on adult category learning until a decade ago focused on classification learning, how people learn to assign items to specified categories. We learn concepts and categories in many ways – such as by interactions, inferences, problem solving, instruction – yet these have received little attention in research on category learning. Not only is much of the laboratory work limited to classification learning, it has rarely varied from a small range of particulars (here is an incomplete list): two categories, small number of features, small number of values per feature, small number of items per category, and divorced from any prior knowledge. Given all the possible ways that even the classification paradigm might be done, these seem very limiting. Of course, it is possible that all the ways of learning and all the possible ways of changing the classification paradigm will not matter in terms of our understanding of category learning, but the evidence suggests just the opposite. As elaborated in the next section, it appears that many of these changes lead to important differences in what is learned.

Second, the items being learned have been limited. In addition to the ways mentioned, almost all have been objects (or descriptions of objects), with little examination of problems, people, situations, scenes, and so on. In addition, the items in most experiments have generally consisted of features only, with no relational structure beyond cooccurrences. None of the main classification theories developed in the exemplar-prototype debates allow relations in their item representations. Given that real-world categories have much relational structure, as well as much prior knowledge, it is unclear how well these findings will relate to more complex cases.

2.29.3 Beyond Classification and Featural Representations

In this section, we consider some recent work that begins to address these limitations. First, research has extended concept and category learning from a focus on classification to consider other means of category learning. Second, we consider two formal models that were designed to examine category learning beyond classification, the Rational Model and SUSTAIN. Third, we review some work that has gone beyond representations of features to ask how more complex categories might be learned, including the influence of prior knowledge. Finally, we consider how far this new research has gone in providing a resolution, or at least a partial resolution, to these limitations.

2.29.3.1 Category Learning Beyond Classification

As we outlined at the beginning, concepts and categories have many functions, of which classification is just one. Classification is an important one: By determining what category an item is in, one has access to much relevant knowledge about that item. However, the near-exclusive focus on classification learning in laboratory experiments is problematic for two reasons. First, if we learn categories in multiple ways, it seems prudent to examine more than classification learning to get a full understanding of category learning. Second, classification learning has an important difference from most other conceptual functions. In classification learning, the goal is to determine what category the item is in. This requires figuring out what distinguishes the competing categories. However, most other functions, such as prediction, understanding, or communication, require using what you know about a particular category, with the other categories often not mattering at all. That is, classification requires distinguishing between categories, but most other functions require within-category knowledge. For example, Chin-Parker and Ross (2004) found that classification learning led to learning only about diagnostic features (those that are predictive of category membership). People learned nothing about the other features that were not predictive of a category, even though they occurred 80% of the time in both categories. This result is exactly as predicted by classification theories, such as the exemplar models mentioned earlier (e.g., Nosofsky, 1986). However, if one is predicting what a new animal eats, it requires knowing more than what type of animal it is; one also needs to know what food is eaten by animals of that type. If one is solving a math problem, knowing the type of problem is helpful only if it allows access to relevant information about how to solve problems of that type.

A number of laboratory tasks have been examined over the last 10 years that extend our understanding of category learning by examining types of learning other than classification. These tasks emphasize how categories allow us to accomplish the goals we have: Predict, solve problems, explain. We focus on three tasks here: Inference learning, in which the classification is provided; category use, in which the learner uses the category to learn some other task; and an unsupervised learning, in which no category information is provided. We give a rough outline of their procedural differences from classification learning in **Table 4** and a rough outline of the processing in **Figure 2**.

First, one can learn about categories by inferring features of category members. Inference learning is a

Table 4 Simplified procedures for various category-learning laboratory tasks

Classification	An item is presented Subject responds with one experimenter-defined category label Feedback given on classification Next item is presented
Inference	An item (one feature missing) is presented, along with category label Subject responds with value of missing feature Feedback given on inference Next item is presented
Category use	An item is presented Subject responds with one experimenter-defined category label Feedback given on classification Subject uses category and item to do some task (inference, problem solve) Feedback given on second task Next item is presented
Unsupervised[a] (using Minda and Ross, 2004, to be specific on procedure)	An item is presented Subject uses item to do some category-related task (prediction) Feedback given on task Next item is presented

[a]Note: no mention made of category, but category is useful for prediction.

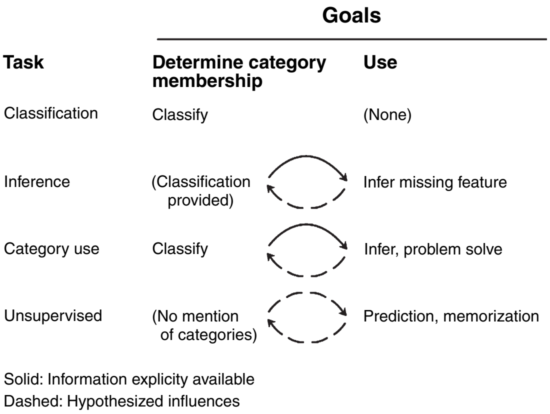

Figure 2 Schematic to illustrate the relation among the different laboratory tasks used to study category learning.

paradigm that can be compared to classification directly: Rather than presenting N features and asking what category the item is in, the label is given along with $N-1$ features, and the person's task is to supply the missing feature (Yamauchi and Markman, 1998). Suppose there is an item in Category A that has three features; we can characterize this item as {A a1 a2 a3}. In a classification trial, the learner would be provided with {? a1 a2 a3} and have to provide the label A, whereas in an inference trial, the learner would get {A a1 a2 ?} and have to provide the missing feature, a3. The procedure is presented in **Table 4**. This task matches a common means of interacting with categories in which one knows the category and parts of the item but has to infer a missing feature: Given this dog, what will it eat? Or given this type of situation, what will happen next? This task appears to be very similar to classification learning but leads to rather different learning. In inference learning, people learn prototypical features that are not diagnostic (Chin-Parker and Ross, 2004) and features that vary in their exact presentation (Yamauchi and Markman, 2000). Focusing within a category leads people to learn what the category is like, not just what distinguishes it from another category (Yamauchi and Markman, 1998; see Markman and Ross, 2003, for a review; Johansen and Kruschke, 2005, for counterarguments). However, if one knows what a category is like (e.g., dogs, permutations), one can often tell if an item is a member of the category as well. In addition, inference learning, by focusing on what the category members are like, may also help people to understand the underlying coherence of relational categories (Erickson et al., 2005).

Second, other work has considered classification as part of a larger goal-related task and asked how category knowledge might be influenced. The rationale is that classification is not usually the goal; usually we want to do something with the classification. We classify an object as a pencil and use it to write, a person as a psychopath and stay away, a math question as a permutations problem and apply the appropriate formula. Ross (1997, 1999, 2000) has conducted research with a category use paradigm in which one not only classifies the item but also then uses the item to perform some goal-oriented task, such as make an inference about the item or solve the problem (see **Table 4** and **Figure 2**). In these studies, the uses of the category (e.g., inference or problem solving) influence the category representation, even for later classification. Performance on later classification tests is not a function just of the diagnosticity of the features but also of their importance for the use. Those features involved in the inference, for example, are viewed as more central to the category than equally predictive features that are not involved in the use. To be more concrete, Ross (1997) had people learn to classify patients (sets of symptoms) into two diseases and then choose the treatment to give that patient (which depended on both the disease and the symptoms). Two symptoms were equally and perfectly predictive of the disease, but only one was also predictive of the treatment. After learning, people were given a single symptom and asked which disease they would think a patient had if this was all they knew of the symptoms. Although both symptoms were perfectly predictive of the disease, learners were much more likely to correctly classify the one that was also predictive of the treatment. Even when we explicitly classify items, we continue to learn about the category from other nonclassification uses, and this knowledge influences a variety of later category-related judgments, including classification. Related results have been found with problem solving (Ross, 1997, 1999), and even with young children (Hayes and Younger, 2004; Ross, et al., 2005).

The goals for which categories are used during learning can have additional consequences that may influence the representation of real-world categories. Brooks et al. (2007) argue that tasks in which classifications are used for a different primary task may lead to a very different learning of the classification knowledge than classification alone. In a clever set of experiments, they show that the other task diverts attention away from analyzing the category structure for classification, such that learners believe the category has defining features when it does not. When we divert our resources to using the category, the knowledge learned from the classification may be a less central part of our category representation.

Third, we can look at cases in which people are not even told there are categories. If we have a helpful teacher/parent, we may get feedback on the category membership (classification) or labeled items (inference), but in many cases, such as much of our informal learning, we may not. In unsupervised learning tasks, the categories are not provided to the learner (e.g., Ahn and Medin, 1992; Heit, 1992; Wattenmaker, 1992; Clapper and Bower, 1994; Billman and Knutson, 1996; Kaplan and Murphy, 1999; Clapper and Bower, 2002; Love, 2003). Here, the particular interactions people have with the items may affect later formation of categories (e.g., Lassaline and Murphy, 1996). Many of these unsupervised tasks focus on the items and the categories, often through observation, memorization, or sorting. When there is a more goal-oriented learning task, such as predicting a critical feature or solving a problem, any learning about the categories or items is incidental to attaining the goal. Minda and Ross (2004) found in a prediction task (where category membership was crucial to the prediction) that the unsupervised learning led to paying more attention to a wider set of features than did classification learning.

In sum, although classification learning has been the dominant paradigm for studying category learning, including other tasks may provide a more complete picture of category learning. Classification is a critical function of categories, but it is critical because it provides access to knowledge about the item that can be used to infer, predict, understand, or explain. Classification learning promotes learning what distinguishes the categories, whereas these other functions of categories tend to promote learning what each category is like.

2.29.3.2 Formal Models That Do More Than Classify: Rational Model and SUSTAIN

Most category-learning models have focused on classification learning, but a few have considered other category functions. We consider two prominent models, J. R. Anderson's (1990, 1991) Rational Model of categorization and SUSTAIN (Love et al., 2004), with an emphasis on learning mechanisms and extensions to multiple functions of categories.

2.29.3.2.1 The Rational Model of categorization

The Rational Model is a component of Anderson's (1990, 1991) rational analysis of cognition. The central claim of the model is that categories serve as a basis for prediction. Classification is a specific type of prediction in which the task is to predict the category label. The model stores clusters of exemplars with similar properties. The primary goal of the model is to develop clusters that optimize the accuracy of a variety of category-based predictions.

Predictions in the model are based on conditional probabilities. The probability that a new exemplar with a set of properties F has property j (on dimension i) is:

$$P_i(j|F) = \sum_k P(k|F)P_i(j|k), \qquad [1]$$

where $P(k|F)$ is the probability that an item belongs to cluster k given that it has properties F, and $P_i(j|k)$ is the probability that an item in cluster k has property j on dimension i. The idea behind equation 1 is that predicting a missing property involves the summing of probabilities across multiple clusters, with the influence of each cluster weighted by its probability, given the new item. The sum of these values across all clusters is the probability that the new item has property j.

Clusters are learned that optimize the accuracy of predictions. The first exemplar forms its own cluster, and each added exemplar is either placed into an existing cluster or a new one that contains only itself, whichever maximizes within-cluster similarities. Thus, a new item that is dissimilar to the earlier items is more likely to begin a new cluster. The probability that a new item will be placed into cluster k depends on (1) how similar the item and cluster features are and (2) the size, or base rate, of the cluster, with the item more likely to be placed into a cluster that represents more items.

Two aspects distinguish this model from the classification models described earlier. First, clusters differ from prototype and exemplar representations in that the goal of clustering is more abstract: to capture statistical structures in the environment suitable for making predictions. Interestingly, clusters sometimes mimic each of the other approaches (cf. Nosofsky, 1991), suggesting that aspects of both prototype and exemplar representations are useful.

Second, the model predicts features as well as category labels, so it can be used for nonclassification tasks, such as inference learning (see Yamauchi and Markman, 1998). In addition, it provides an explanation for how people might induce a missing feature of an exemplar when the category is unknown. Suppose that you hear an animal rustling behind a bush, and you think it is probably a dog but possibly a raccoon. What is the probability that the animal barks? In terms of equation 1, dog and raccoon are different values of *k*. Murphy and Ross (1994, 2005; Malt et al., 1995; Ross and Murphy, 1996) tested this hypothesis in a large number of studies and usually found evidence against it. Their work shows that people instead tend to base predictions on the most likely category only (dog in the previous example). Although the Rational Model was not supported, it provided a strong alternative view and has led to a consideration of when single categories and multiple categories might be used for predictions.

2.29.3.2.2 SUSTAIN

SUSTAIN (Supervised and Unsupervised STratified Adaptive Incremental Network; Love et al., 2004) is a network model of category learning that shows great flexibility compared with prototype and exemplar models for two main reasons. First, like the Rational Model, it seeks to build clusters that capture regularities or structure in the environment. Second, unlike the Rational Model, this search for structure is guided by the goals of the learner. This goal-oriented learning allows the model to account for a wide variety of category learning results beyond classification.

Categories in SUSTAIN are represented by clusters. Unlike the Rational Model, these clusters are not collections of exemplars but, rather, summary representations of encountered exemplars. The clusters formed are influenced by the goals of the learner, such as increased attention to the feature (including category label) being predicted and the features most relevant for this prediction. The details of the model are complex, so we outline the main steps of performance and learning here, then turn to a simple example.

The performance, such as prediction of a feature, relies upon the clusters. The item is compared to each cluster, and the most similar cluster determines the prediction. The summary representation of each cluster includes a distribution of expectations for each dimension (e.g., how likely the different values are to occur). The item's features are compared to these various distributions, with selective attention occurring through the tuning of receptive fields on each dimension (akin to visual receptive fields that are, for example, sensitive to a small range of orientations at particular locations). The activation of the cluster increases with the similarity of the item's features to the summary representation (weighted by the selective attention weights). The different clusters can be thought of as trying to explain the input for the particular goal, with lateral inhibition among the clusters leading to a winning cluster that determines the output. Thus, unlike the Rational Model that sums over all the clusters (see equation 1), only the most likely cluster is used to determine the prediction (consistent with the Murphy and Ross, 1994, results).

SUSTAIN is biased toward simple category representations (i.e., few clusters) but does develop more elaborate clustering schemes for complex stimulus sets. Learning involves updating old clusters and developing new ones. In supervised learning, if the correct prediction is made, SUSTAIN will compare the output of the winning cluster to the target response and make small adjustments in receptive field tunings and summary representation values in the direction of the target values, if needed. These changes will lead to a repetition of the item producing a cluster output closer to the target values. If an incorrect prediction is made, a new cluster is created that is centered around the item. (In unsupervised learning, a new cluster is formed if the current item is not sufficiently similar to any existing cluster.)

This explanation is a bit abstract, so we illustrate with a simple example. Imagine there is an object that can be described by three binary features: shading (filled = 0, empty = 1), color (blue = 0, red = 1), and shape (circle = 0, triangle = 1). Thus, we can represent an (empty red circle) as (1, 1, 0). Suppose this was the first item presented, then Cluster 1 (CL1a) would simply be (1, 1, 0) as seen in **Figure 3**. Now suppose the second item was (empty red triangle), item (1, 1, 1), and it was similar enough to be put in the same cluster. The updated Cluster 1 (CL1b) would be adjusted to be (1, 1, 0.5). The last value does not represent a triangular circle but, rather, an increased probability that a new item represented by that cluster will be a triangle. If the third stimulus (empty blue circle), (1, 0, 0), is not sufficiently similar to CL1b (unsupervised learning) or does not predict the correct response (supervised learning), the model recruits a new cluster CL2a to represent that item, as shown in the figure.

Recall that SUSTAIN prefers simple cluster sets when possible. To demonstrate, if the stimuli in **Figure 3** were divided by shape into two categories, it is likely that SUSTAIN would develop two clusters, each at the center of the front and back face of the cube, (0.5, 0.5, 0) and (0.5, 0.5, 1). These would indicate a high probability of circle//triangle for the front/back cluster and intermediate probabilities for other dimensions. This is a simple clustering solution, because it strongly emphasizes just one feature.

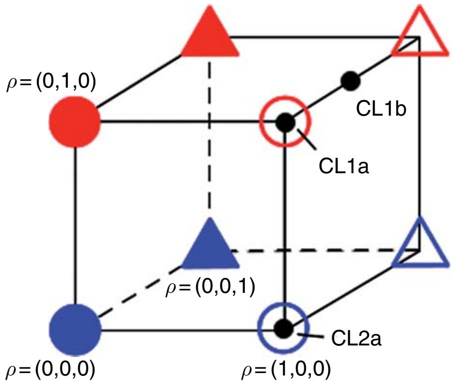

Figure 3 Binary-valued stimuli to illustrate working of SUSTAIN. Coordinates are shading (0 = filled, 1 = empty), color (0 = blue, 1 = red), and shape (0 = circle, 1 = triangle). Black dots represent clusters CL1a, CL1b, and CL2a following the presentation of items (1, 1, 0), (1, 1, 1), and (1, 0, 0), respectively.

Like the Rational Model, the clusters in SUSTAIN adapt to their learning environment, but they are also sensitive to the goals of the learner (e.g., Barsalou, 1985; Solomon et al., 1999; Medin, et al., 2006). In supervised learning, new clusters are created in response to an incorrect decision, allowing the model to adjust to specific learning criteria. This differs from clustering in the Rational Model, where new clusters are created in response to exemplars that are generally dissimilar to existing clusters.

SUSTAIN's ability to adapt to different category functions is also seen in comparing classification and inference learning. Using a family resemblance structure (where all items are similar to a prototype for the category), Love et al. (2000) found that inference learning often led to a single cluster (i.e., close to the prototype), whereas classification tended to create several clusters per category, indicating the use of simple rules and memorization when classifying. (See also Love and Gureckis, 2005, for a related application to a real-world difference in goals.)

In summary, SUSTAIN focuses on the relationship between learning goals and category structure. In a sense, SUSTAIN generalizes the contributions of the Rational Model by proposing that, in addition to capturing the structure of the environment, categories also capture a learner's goal-specified relationship with the environment.

2.29.3.3 Beyond Featural Representations

Much of the current work on category learning is limited not just in how the categories are learned but also in what is learned. The categories learned in these studies are different from real-world categories in two important ways. First, the category structure is feature-based with only very simple relations among the features. Most, if not all, real-world categories have much more complex structures, and relations are essential. Second, the materials are usually devised to make as little contact with any world knowledge as possible, to allow an unconfounded examination of category learning. However, the learning of most (and perhaps all) real-world categories is influenced by our world knowledge: By minimizing the influence, we may be ignoring a major influence in the learning. We discuss each of these ideas with a brief examination of relevant models.

2.29.3.3.1 Relational information

Most of the categories studied in experimental settings with adults consist of a small number of features (usually 2–5), with some of the features predictive of each category. These simple structures, cleverly designed, are often sufficient to provide tests of particular aspects of current formal models. The underlying (implicit) assumption is either that real-world categories have such structures (which they do not) or that the learning principles derived from studying such simple structures will apply to learning more complex structures. This latter possibility remains feasible, though we present evidence throughout this section suggesting that there are many important differences. The main point we wish to make here is that the category structures that have typically been examined are a very small subset of the possible structures and do not have much to do with real-world category structures (see Murphy, 2005).

Let's take a simple example of relational structure, within-category correlations. We can classify an item as a bike using features such as handlebars, tires, and such. Bicycle handlebars include ones that are either dropped or straight, and bicycle tires include ones that are knobby or slick. However, these are not independent features: Dropped handlebars and slick tires usually go together (racing bikes), as do straight handlebars and knobby tires (mountain bikes). These within-category correlations are common in many real-world categories: They do not add to the category predictiveness of the features, but they are important as signals to additional category structure, such as subcategories. When within-category correlations have been examined in laboratory studies of classification learning, they do not appear to be learned easily, if at all (Chin-Parker and Ross, 2002). This difficulty is generally consistent with classification learning models, because the within-category correlations do not improve the classification predictiveness. Despite this unanimity of classification models and laboratory classification learning data, people do learn within-category correlations in the world. They even learn it in the laboratory when given inference learning (Chin-Parker and Ross, 2002) or when prior knowledge promotes the correlation (Murphy and Wisniewski, 1989).

The point is that as one moves to more complex relations among features, we know little about category learning. Of course, the structure of real-world categories is far more complex than a single within-category correlation: Imagine the structure for carrots or permutations. The dominant classification models are all feature based and do not allow for complex relational information (though see Pazzani, 1991). The problem is that adding relational information greatly complicates models and brings a host of additional issues that need to be addressed (e.g., Hummel and Holyoak, 2003).

2.29.3.3.2 Knowledge

What people know affects what they learn: Cricket fans gain much more from watching a cricket match than do those of us who know nothing about the game. Not surprisingly, the effect of knowledge is also true in category learning: Knowledge can have large effects on the ease of learning as well as what is learned (see Murphy, 2002: chapter 6, for a review). Continuing the example on within-category correlations, knowledge has a large effect. Murphy and Wisniewski (1989) had people learn categories containing within-category correlations. For example, their materials included correlated features that were conceptually related (for a clothing category: worn in winter, made of heavy material) or not (blue, machine washable). They found that in the absence of knowledge relating the features, classification learners did not learn these within-category correlations. However, in the condition in which knowledge related the features, classification learners did show sensitivity to the correlations. Ahn and her colleagues argue that there are so many possible correlations in the world that people cannot notice all of them, so people use their knowledge to notice correlations that are meaningful to them (e.g., Ahn et al., 2002; though see McRae, 2004).

We mention briefly three other effects of knowledge on learning. First, the learning of new categories that are consistent with prior knowledge is greatly accelerated compared with learning unrelated or inconsistent categories. For example, Wattenmaker et al. (1986) showed that categories consisting of items whose features related to a theme such as honesty (e.g., returned the wallet he had found in the park) were faster to learn than categories whose items had unrelated features. Second, even learning to classify items into categories consistent with knowledge but for which a prior concept is unlikely to be available (e.g., arctic vehicles) is faster than learning to classify into unrelated categories (e.g., Murphy and Allopenna, 1994). Third, the learning influence of prior knowledge does not restrict learning to only those features related to the prior knowledge, but generally, those are learned more quickly than the unrelated features of the items (e.g., Kaplan and Murphy, 2000).

How does knowledge influence category learning? Wisniewski and Medin (1994) suggest three important possibilities (also see Heit and Bott, 2000). First, knowledge might weight or select features of the item. Second, knowledge might allow one to infer additional, relevant features. Third, although these first two possibilities view knowledge as independent of the learning process, knowledge and learning may be more tightly coupled or interactive. Learning may influence the activated knowledge, which may then influence later learning. As one example, learners might begin to change how they interpret some features as they see how the features relate to their goal. The main point of this work is that knowledge and concept/category learning cannot be thought of separately: Our knowledge of particular concepts is intimately intertwined with other knowledge, and we use that other knowledge both to help learn new information about concepts and that this learning in turn may influence our other knowledge.

Despite this large influence of knowledge on category learning, most category learning research has examined cases in which knowledge influences are minimized (Murphy, 2005). If there are interactions between the influences of knowledge and the learning, some (unknown) part of what we learn from examining category learning in the absence of knowledge may not be applicable to the cases in which knowledge is used. In addition, there may be learning processes with knowledge that are not required in the knowledge-free classification learning experiments.

There has been some progress in considering how to account for the influence of prior knowledge on category learning in a more general way, including

Heit's Baywatch model (Bayesian and empirical learning model, e.g., Heit and Bott, 2000) and KRES (Knowledge-Resonance model; Rehder and Murphy, 2003). Space precludes much description of these complex connectionist models, but both models add nodes to represent prior knowledge in addition to the usual ones to represent the features of the items. The presented item activates its features but also activates prior knowledge that it is related to, providing an additional source of activation for the category decision. The models account for a wide variety of data. For example, KRES predicts that the learning of categories consistent with prior knowledge is accelerated compared with unrelated categories, and even features unrelated to the prior knowledge are learned. In addition, this model incorporates the interactive view between knowledge and learning by the learning influencing the connection weights between features and between features and prior knowledge.

These models of prior knowledge do not include relational representations. Relational models are beginning to be developed (e.g., Hummel and Ross, 2006; Kemp et al., 2006; Tomlinson and Love, 2006) but need also to address the pervasive influences of prior knowledge.

2.29.3.4 Directions for Providing Integration

We have considered some limitations of the current category learning work both in terms of the learning tasks and in terms of the featural representations. Different means of learning about categories provide a variety of knowledge about the category not just for classification but also to support all the category-based cognitive activities. In addition, the knowledge about a category has to be intimately related to our other conceptual knowledge to be useful. What does this suggest about concept and category learning?

A main lesson that we have taken from a consideration of these various limitations is that the study of concepts and category learning needs to be integrated into other areas of cognition. Conceptual knowledge needs to support many cognitive activities, not just classification, and examining these category-based activities across a variety of domains will both point out places in which we need to further our understanding of category learning and help to make the work on category learning more relevant to other areas of cognition. Much of our learning depends on our goals, so considering more than category learning is an important part of ensuring this integration.

Murphy and Medin (1985), in a seminal paper, proposed that the study of conceptual structure could not rely on similarity to explain why objects might cohere, or go together, in a category. Although similarity might be a useful heuristic in some cases, it is too unconstrained to provide an explanation of category coherence. They proposed that the coherence of the category depended on its fit to people's prior knowledge – the naive theories people have. This proposal changes the idea that the category members are similar to that they have some similar underlying rationale. Their internal structure is defined not just by features but also by relations connecting features. In addition, their external relations are critical: They must relate somewhat consistently with other knowledge the person has. This view is often called the theory view to make clear that it views category coherence as depending upon people's theories, not simple similarity.

This proposal has had a major influence on how conceptual structure is thought about and investigated. It was instrumental in leading to much of the work on how knowledge influences category learning. We mention two interesting illustrations. First, Wisniewski and Medin (1994) gave subjects a set of children's drawings of people and asked them to provide a rule for each category; they were told either that one group of drawings was from creative children and the other from noncreative children or, for other subjects, that one group was from city children and the other from farm children. The rules generated were very different, picking up on aspects of the drawings that were consistent with some ideas of those types of people. For example, one feature was seen as a pocket, indicating detail, when it was from the drawings by creative children, but was seen as a purse when it was from drawings of city children. Second, Ahn and Kim (2005; Kim and Ahn, 2002) have examined clinical psychologists' understanding of various psychodiagnostic categories (such as depression or anorexia). Although the training they receive emphasizes a prototype representation (classification is often in terms of some criterial number of features being present), the clinical psychologists often apply their causal theories of the disorders to help diagnose and determine treatments. Thus, this theory view has led to a wealth of interesting research relating prior knowledge to concept learning. The main shortcoming of this view is the lack of specific details on how knowledge influences learning.

There are likely to be many cases in the learning of complex categories in which the constraints imposed by theories are not sufficient even for classification. For example, in learning to classify members of many real-world categories, there may be hundreds of potential features that could be important for determining category membership, so what determines which features people use? Knowledge may reduce the number of likely features and relations but still leave too large a number to consider. One possibility is that the importance of the features and relations for the overall goal of the task may provide a heuristic as to which ones are important for classification. For example, one might not be able to tell how to classify a particular math problem, but as one gets experience solving problems of that type, those aspects of the problem critical for solution may provide a good clue as to how to classify future problems. The comparison to other category members with respect to these useful features may also help to lead to a deeper understanding of why the problem is solved in this way.

It is important to clarify this suggestion and make clear how it relates to the general integration goal of this chapter. Concepts and categories support a variety of functions. The usefulness of this knowledge across the different functions provides constraints that one cannot get from a single function. In addition, the changes to the representations as one both uses and gets feedback on one function provides knowledge that can be used for other functions. For example, learning to classify complex items, with many features and relations, is a very difficult task if one relies only on feedback from the classification (which may be why classification learning experiments typically use few features and values). However, if these same categories are used for inferences or problem solving, that use provides suggestions as to what features and relations one might consider. Similarly, background knowledge can be used to help focus on relevant features and relations or even to learn new features that are important for later classifications (e.g., Wisniewski and Medin, 1994). The apparent difficulty of learning complex concepts is partly a result of thinking about it as some isolable process that relies only on classification and feedback on the classification. People have many sources of information from both the various interactions with the items and their prior knowledge to help in learning new concepts and categories.

2.29.4 Integrating Concepts and Categories into Cognition

We have been arguing throughout this chapter that it is important to think about concept and category, learning more broadly to integrate it into the many cognitive activities in which they play such a critical role. In this section, we illustrate this possibility by examining concepts and categories in two very different areas, problem solving and language.

2.29.4.1 Problem Solving

Categories play a critical role in human problem solving. Being able to correctly classify a problem as a permutations problem allows you to recall and apply the appropriate formula. In this section, we describe how category knowledge can influence various aspects of the problem-solving process, how problem categories change with experience, and how problem solving can affect the category representation.

Most models of problem solving consist of some version of the following five stages: (1) problem identification and creating a mental representation of the initial problem state and goal; (2) identifying and selecting a set of operators, procedures, or strategies to make progress toward that goal; (3) applying those operators and generating a solution; (4) assessing whether the solution satisfies the goal; and (5) storing the solution with other knowledge about the problem/category (Newell and Simon, 1972; Bransford and Stein, 1993; Pretz et al., 2003).

Categories impact all aspects of this process and are especially critical in the early stages (**Figure 4**). The process of problem identification is a classification that determines whether or not the current problem is like other problems encountered in the past. After the problem is classified, the problem solver can then recall and apply a set of procedures, strategies, or rules to solve the problem, such as recalling the appropriate formula for a permutations problem. Category knowledge may also be helpful in later stages of problem solving, such as evaluating whether or not a potential solution satisfies the known constraints of the problem type.

The problem goal is also important. Since category knowledge is used in the service of accomplishing some particular task, knowing how the goal relates to the problem features is a critical part of understanding the problem and identifying the appropriate solution procedures to solve it. As an illustration,

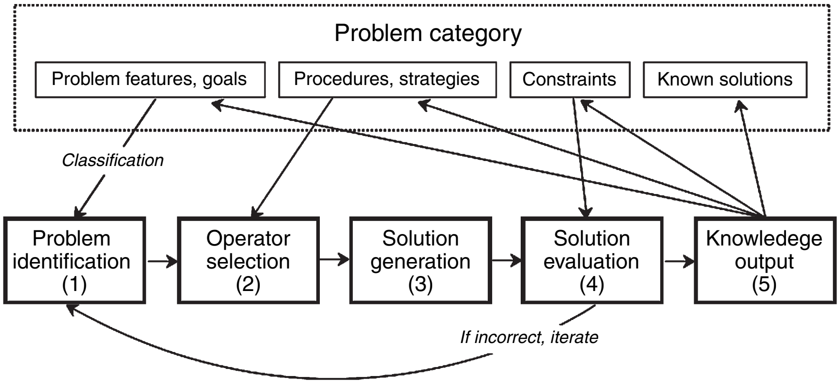

Figure 4 Five stages of problem solving (lower boxes) and relation to knowledge of problem categories (upper box). Arrows connecting boxes in the lower portion of the figure represent the general order of processing. Arrows connecting boxes between the upper and lower portions represent interactions between problem solving and category knowledge.

imagine that, as you are ready to drive to work, you notice the lights were left on, and the car does not start. If you classify the situation as a dead car battery, you can search your relevant category knowledge, perhaps recalling that one solution is to use jumper cables and ask your next-door neighbor for a jump. However, if your primary goal was to get to work on time, you might access other relevant knowledge such as calling a friend for a ride. For categories to be helpful for problem solving, they must go beyond simple diagnostic features and include information about the problem or goal procedures or strategies to accomplish that goal, and knowledge to evaluate the solution.

2.29.4.1.1 Differences between expert and novice category representations

Much research has shown that experts categorize problems with the principle or structural features of the domain, whereas novices rely on using the surface features (Larkin et al., 1980; Bedard and Chi, 1992). Chi et al. (1981) found that when physics experts were asked to sort a set of problems into those that could be solved similarly, they sorted on the basis of the underlying physics principles, such as Newton's Second Law, whereas novices sorted on the basis of the surface characteristics, such as problems with inclined planes. The physics experts had learned to associate the deep principles with the problem features and could take advantage of this knowledge for problem categorization.

The finding that experts use the deep principles of the domain to classify and reason about novel problems has been shown in many other domains including chess (Chase and Simon, 1973; Gobet and Simon, 2000), computer programming (Adelson, 1981), electronics (Egan and Schwartz, 1979), and mathematics (Silver, 1979, 1981), among others. How does one go from a novice category representation to that of an expert? One suggestion from the problem-solving literature is that much of the learning comes as a by-product of the problem-solving activity itself (Ross and Kennedy, 1990; Cummins, 1992). There are multiple possibilities for how this can occur, including adding new knowledge to the category, modifying previous knowledge by weighting particular category features, adding constraints to the category to further specify the category boundary, deleting inappropriate knowledge, or acquiring new, more specific categories. These possibilities suggest not only that categories are a critical aspect of human problem solving but that category representations can be adapted and changed through problem solving.

2.29.4.1.2 Problem solving and category learning

Much experimental work shows that problem solving can affect category learning (Ross, 1996, 1997, 1999). For example, Ross (1996, Experiment 2a) conducted a category use experiment (similar to the ones mentioned in the section 'Beyond classification and featural representations') in which students learned to classify equations into two categories and then solve them. The solutions of the equations differed, and the question was whether this would affect the participants' category representations and influence their future classifications. Table 5 shows a sample of the materials. Every problem had an x and a y variable, and half of the participants solved for variable x and the other half for variable y. The equations were

Table 5 Sample materials from Ross (1996)

Phase of experiment	*Equation*	*Solution*	*Solve for* x	*Solve for* y
Study				
Type 1	$9 + \frac{sx}{t} = \frac{5qy + b}{f}$		SMD	MSD
Type 2	$sp + \frac{ry}{2} = \frac{g + 9x}{b}$		MSD	SMD
Test				
	$k + \frac{9z}{a} = \frac{s + r}{t}$	SMD	Type 1	Type 2
	$s + \frac{2}{n} = \frac{9 + mz}{p}$	MSD	Type 2	Type 1

SMD, subtract, multiply, divide; MSD, multiply, subtract, divide. At test all subjects solved for *z*. The ''solve for *x*'' and ''solve for *y*'' labels refer to the study conditions.

From Ross BH (1996) Category representations and the effects of interacting with instances. *J. Exp. Psychol. Learn. Mem. Cogn*. 22: 1249–1265. Copyright 1996 by the American Psychological Association. Reprinted with permission of the author.

created so that those who solved for variable *x* would use a subtract-multiply-divide (SMD) solution procedure on type 1 problems and a multiply-subtract-divide (MSD) solution procedure on type 2 problems, whereas those who solved for variable *y* would use the opposite procedures on the two problem types. At test, they classified novel test problems and solved for a new variable *z*. Tests requiring a SMD solution procedure tended to be classified as type 1 problems for those participants who solved for *x* and as type 2 problems for those who solved for *y* (and vice-versa for test problems that required the MSD solution procedure). Although the two groups classified the same problems into the same categories during learning, they classified the test items into opposite categories because of how they interacted with the items. The way in which problems are solved can influence which categories people have.

Beyond traditional problem solving, Medin and colleagues show strong influences of extended interactions on how tree experts classify and reason (Medin et al., 1997; Lynch et al., 2000; Proffitt et al., 2000). Taxonomists and maintenance workers sorted a set of trees on the basis of morphological features, but each weighted the importance of those features differently, whereas the landscapers sorted more on the basis of utilitarian features, such as providing shade or ornamental quality (Medin et al., 1997). The experts' category representations were influenced by how they interacted with items in the category. Proffitt et al. (2000) found that all of the groups used ecological-causal domain knowledge in addition to general taxonomic knowledge to make inductions. These results are consistent with the laboratory results showing that experience interacting in a domain influences a person's category representation and subsequent reasoning from those categories.

The role of concepts and categories in problem solving is pervasive. For categories to be useful for problem solving, they require more than simple featural representations, including information about the goals, solution procedures, and constraints of the problem. Problem categories change with experience: There are general shifts from surface feature representations to structural (relational) representations, as well as influences of the particular uses.

2.29.4.2 Language

Categories are also critical in language use. We focus on one area within language performance (i.e., the encoding of syntactic number) to explore different kinds of categories, the processes that operate on them, and the functions such categories serve in language performance.

In general, the types of categorical structures necessary to support language processing seem to depart considerably from those focused on in category learning research. For example, during language production, transforming a message that is full of meaning into an utterance that is full of sound (Bock and Miller, 1991) may involve (at least) coordinating systems of categories corresponding to syntactic structure (e.g., the syntax of phrases or sentences), grammatical functions and thematic roles (e.g., subject, object, agent, patient), word types (e.g., noun, verb), word-specific grammatical information (e.g., grammatical gender and number), word meanings, morphophonology, and prosodic structure. Particular to syntactic number, the production of a lexical singular (e.g., cat/argument/dustbuster) or a lexical plural (e.g., cats/arguments/dustbusters) is thought to be rooted in a kind of categorization involving the apprehension of the referent of a noun as a single thing or more than one thing (Eberhard et al., 2005). Some categorizations are not simple: A bowl full of "fresh or dried food that is usually made from flour, eggs, and water formed into a variety of shapes" may be conceived of as a mass or as comprising individual units and so the speaker may elect to call it pasta or noodles, respectively. This categorization determines the appropriate lexical number characteristics of the noun.

Consistent with our emphasis on category use, assigning lexical number (singular or plural) is not an end in and of itself – lexical number characteristics serve the function of communicating to a listener the numerosity of referents, and they are critical in the computation of grammatical agreement (such as between a subject and a verb; Eberhard et al., 2005; but see Vigliocco et al., 1996). Grammatical agreement in turn serves important communicative functions such as linking pronouns to their referents and helping listeners syntactically bind subjects to their predicates when syntactic ambiguity arises. For example, subject–verb agreement helps disambiguate who has rabies in "The *dog* chasing the *men* that *has* rabies."

Beyond categorizing a referent as one thing or more than one thing, singular nouns in many languages divide into count nouns and mass nouns. The count/mass distinction is thought by some to reflect distinct modes of construal relevant to individuation and allows an interesting examination of concepts in language use.

Count nouns like animal(s), argument(s), and noodle(s) must have a determiner in the singular form (*Animal is fierce), are regularly pluralized, and take the quantifiers many and few. Mass nouns such as wildlife, evidence, and pasta do not need to take a determiner in the singular (Wildlife is flourishing), are not regularly pluralized, and take the quantifiers much or little. What are the psychological implications of this distinction? The cognitive individuation hypothesis proposes that count nouns denote individuated entities and mass nouns denote nonindividuated entities (Mufwene, 1984; Wierzbicka, 1988; Jackendoff, 1991; Bloom, 1994; Bloom and Kelemen, 1995; Bloom, 1996; Wisniewski et al., 1996; Wisniewski et al., 2003; Middleton et al., 2004). The class of individuated entities includes common objects such as cats, blenders, and airplanes but also includes things bounded spatially (even to an absence of matter, e.g., a hole, Giralt and Bloom, 2000) or temporally (events, such as a footrace or a party; Bloom, 1990). Individuation can apply to entities linked by common fate or goal (e.g., a gang, a flock) or common purpose (a bikini may be conceived as an individual because the two pieces perform one function; Bloom, 1996), as well as from a variety of other factors (see Goldmeier, 1972; Jackendoff, 1991; Soja et al., 1991; Bloom, 1994).

2.29.4.2.1 Categorization and cognitive individuation

The process of individuation is not just a categorization based on the physical features on an entity. Cognitive individuation involves active construal of an entity, which can be flexibly applied and has important consequences. Specifically, if a person individuates an entity, that person predicates that features of the entity must hold specific functional relationships to each other (Wisniewski et al., 2003; Middleton et al., 2004). For example, if one individuates a configuration of wood as a table, one is comprehending how the configuration of four upright pieces of wood and a horizontal wooden plane go together to support the important function of supporting stuff. This construal does not allow pieces to be randomly removed or rearranged. In contrast, if one categorizes the table as a nonindividuated entity, one might focus on the material rather than the configuration. If so, one might predicate the important property of 'is flammable,' which does not depend on the configuration.

Evidence that individuation is a flexible process in which different outcomes (e.g., individuation vs. nonindividuation) can arise given the same stimulus was reported by Middleton et al. (2004; Experiment 4). One group of participants viewed a bounded pile of coarse decorative sugar in a box (a novel stimulus) and chose to refer to it with count or mass syntax. A second group viewed the stimulus, followed by a mode of singular interaction where they repeatedly took an individual grain and placed it through one of several holes in a rectangular piece of cardboard. Participants in this second group were more likely to refer to the stimulus with count syntax than the control group. This demonstrates that individuation is not directly tied to the features of a stimulus. Rather in this case, how a person interacted with the entity was related to its individuation status, and this in turn was reflected in the syntax they used. This point introduces the functionality of the count/mass distinction: Using count or mass syntax provides a means to communicate distinct construals of an entity in terms of individuation. This may be particularly useful when the mode of construal as an individual or nonindividuated entity is atypical for a referent. Consider 'too much woman' as in "[S]he's [Jennifer Lopez] too much woman for that piece of snore [Ben Affleck]." Using mass syntax with what is typically a count noun (i.e., woman) allows the speaker to communicate construal of womanly attributes as lying on a continuum, with Jennnifer Lopez falling on the high end (at least, too high for Ben Affleck). Attributes of other common objects can be construed as lying on a continuum, as communicated in "[M]any border collies are destroyed because they proved to be too much dog for their owners," where 'dogness' may be some value along a continuum composed of activity level, obedience, ferocity, and so on. (These examples are extracted from American Web sites.)

Language may not just reflect concepts, it may influence the representation of concepts. The boundaries of basic categories may not be invulnerable to the effects of language (e.g., Boroditsky, 2001; Gordon, 2004). As one example, Imai and Gentner (1997; see also Imai, 1995) showed that Japanese- and English-speaking children differentially weighted the importance of a similar substance and similar configuration when choosing which item was the same as an example. This issue of how language may lead to differences in categorical structure is potentially very important inasmuch as we learn a large proportion of our concepts through communicating by direct instruction or implicitly through conversation (see Markman and Makin, 1998).

2.29.5 Conclusions

In this chapter, we have presented an overview of research on concept and category learning, but we have done so from a particular perspective. Although classification models and experiments have dominated the laboratory work in this area, some recent work has questioned both the focus on classification and its use of simple featural-based items. This work has promoted a broader examination of concept and category learning in three ways. First, a variety of category-learning paradigms are being investigated, along with models that can perform other category functions besides classification. Second, the complexity of the material being learned has increased to include relational categories, prior knowledge, and nonobject categories. New models are also being proposed to begin to account for these complexities. Third, this perspective encourages a consideration of how concept and category learning may be viewed in other areas of cognition. These advances should provide a richer, broader view in the future, so we can better understand the learning of concepts and categories and their crucial role for intelligent thought and action.

References

Adelson B (1981) Problem solving and the development of abstract categories in programming languages. *Mem. Cognit.* 9: 422–433.

Ahn W (1999) Effect of causal structure on category construction. *Mem. Cognit.* 27: 1008–1023.

Ahn W and Kim NS (2005) The effect of causal theories on mental disorder diagnosis. In: Ahn W, Goldstone R, Love B, Markman A, and Wolff P (eds.) *Categorization Inside and Outside the Laboratory: Festschrift in Honor of Douglas L Medin*, pp. 273–288. Washington, DC: American Psychological Association.

Ahn W and Medin DL (1992) A two-stage model of category construction. *Cognit. Sci.* 16: 81–122.

Ahn W, Marsh JK, Luhmann CC, and Lee K (2002) Effect of theory-based feature correlations on typicality judgments. *Mem. Cognit.* 30: 107–118.

Anderson AL, Ross BH, and Chin-Parker S (2002) A further investigation of category learning by inference. *Mem. Cognit.* 30: 119–128.

Anderson JR (1990) *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.

Anderson JR (1991) The adaptive nature of human categorization. *Psychol. Rev.* 98: 409–429.

Ashby FG and Ennis JM (2006) The role of the basal ganglia in category learning. In: Ross BH (ed.) *The Psychology of Learning and Motivation,* vol. 46, pp. 1–36. San Diego: Academic Press.

Barsalou LW (1985) Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 11: 629–654.

Barsalou LW (1990) On the indistinguishability of exemplar memory and abstraction in category representation. In: Srull TK and Wyer RS (eds.) *Advances in Social Cognition, Volume III: Content and Process Specificity in the Effects of Prior Experiences*, pp. 61–88. Hillsdale, NJ: Erlbaum.

Bedard J and Chi MTH (1992) Expertise. *Curr. Direct. Psychol. Sci.* 1: 135–139.

Billman D and Knuston JF (1996) Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *J. Exp. Pyschol. Learn. Mem. Cognit.* 22: 458–475.

Bloom P (1990) *Semantic Structure and Language Development.* PhD Thesis, Massachusetts Institute of Technology.

Bloom P (1994) Possible names: The role of syntax–semantics mappings in the acquisition of nominals. *Lingua* 92: 297–329.

Bloom P (1996) Possible individuals in language and cognition. *Psychol. Sci.* 5: 90–94.

Bloom P and Kelemen D (1995) Syntactic cues in the acquisition of collective nouns. *Cognition* 56: 1–30.

Bock K and Miller CA (1991) Broken agreement. *Cognit. Psychol.* 23: 45–93.

Boroditsky L (2001) Does language shape thought? Mandarin and English speakers' conceptions of time. *Cognit. Psychol.* 43: 1–22.

Bransford JD and Stein BS (1993) *The Ideal Problem Solver*. New York: WH Freeman.

Brooks LR, Squire-Graydon R, and Wood TJ (2007) Diversion of attention in everyday concept learning: Identification in the service of use. *Mem. Cognit.* 35: 1–14.

Bruner JS, Goodnow JJ, and Austin GA (1956) *A Study of Thinking*. New York: Wiley.

Chase WG and Simon HA (1973) Perception in chess. *Cognit. Psychol.* 4: 55–81.

Chi MTH, Feltovitch P, and Glaser R (1981) Categorization and representation of physics problems by experts and novices. *Cognit. Sci.* 5: 121–152.

Chin-Parker S and Ross BH (2002) The effect of category learning on sensitivity to within-category correlations. *Mem. Cognit.* 30: 353–362.

Chin-Parker S and Ross BH (2004) Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 30: 216–226.

Clapper JP and Bower GH (1994) Category invention in unsupervised learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 20: 443–460.

Clapper JP and Bower GH (2002) Adaptive categorization in unsupervised learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 28: 908–923.

Cummins DD (1992) Role of analogical reasoning in the induction of problem categories. *J. Exp. Psychol. Learn. Mem. Cognit.* 18: 1103–1124.

Eberhard KM, Cutting JC, and Bock K (2005) Making syntax of sense: Number agreement in sentence production. *Psychol. Rev.* 112: 531–559.

Egan DE and Schwartz BJ (1979) Chunking in recall of symbolic drawings. *Mem. Cognit.* 7: 149–159.

Erickson MA and Kruschke JK (1998) Rules and exemplars in category learning. *J. Exp. Psychol. Gen.* 127: 107–140.

Erickson JE, Chin-Parker S, and Ross BH (2005) Inference and classification learning of abstract coherent categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 31: 86–99.

Giralt N and Bloom P (2000) How special are objects? Children's reasoning about objects, parts, and holes. *Psychol. Sci.* 11: 497–501.

Gobet F and Simon HA (2000) Five seconds or sixty? Presentation time in expert memory. *Cognit. Sci.* 24: 651–682.

Goldmeier E (1972) Similarity in visually perceived forms. *Psychol. Issues* 8: 14–135.

Gordon P (2004) Numerical cognition without words: Evidence from Amazonia. *Science* 306: 496–499.

Hampton JA (1993) Prototype models of concept representation. In: Van Mechelen I, Hampton JA, Michalski RS, and Theuns P (eds.) *Categories and Concepts: Theoretical Views and Inductive Data Analysis*, pp. 67–95. San Diego: Academic Press.

Hampton JA (2006) Concepts as prototypes. In: Ross BH (ed.) *The Psychology of Learning and Motivation,* vol. 46, pp. 79–113. San Diego: Academic Press.

Hayes BK and Younger K (2004) Category-use effects in children. *Child Devel.* 75: 1719–1732.

Heit E (1992) Categorization using chains of examples. *Cognit. Psychol.* 24: 341–380.

Heit E and Bott L (2000) Knowledge selection in category learning. In: Medlin DL (ed.) *Psychology of Learning and Motivation,* vol. 39, pp. 163–199. San Diego: Academic Press.

Hummel JE and Holyoak KJ (2003) A symbolic-connectionist theory of relational inference and generalization. *Psychol. Rev.* 110: 220–264.

Hummel JE and Ross BH (2006) Relating category coherence and analogy: Simulating category use with a model of relational reasoning. In: *Proceedings of the Twenty-Eighth Annual Cognitive Science Society*, p. 2647. Hillsdale, NJ: Erlbaum.

Imai M (1995) Constraint on word-learning constraints. *Jpn. Psychol. Res.* 41: 5–20.

Imai M and Gentner D (1997) A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. *Cognition* 62: 169–200.

Jackendoff R (1991) Parts and boundaries. *Cognition* 78: B17–B25.

Johansen MK and Kruschke JK (2005) Category representation for classification and feature inference. *J Exp. Psychol Learn. Mem. Cognit.* 31: 1433–1458.

Kaplan AS and Murphy GL (1999) The acquisition of category structure in unsupervised learning. *Mem. Cognit.* 27: 699–712.

Kaplan AS and Murphy GL (2000) Category learning with minimal prior knowledge. *J. Exp. Psychol. Learn. Mem. Cogn.* 26: 829–846.

Kemp C, Tenenbaum JB, Griffiths TL, Yamada T, and Ueda N (2006) Learning systems of concepts with an infinite relational model. Twenty-Fourth National Conference on Artificial Intelligence (AAAI-06), Boston, MA, July 2006.

Kim N and Ahn W (2002) Clinical psychologists' theory-based representations of mental disorders predict their diagnostic reasoning and memory. *J. Exp. Psychol. Gen.* 131: 451–476.

Kruschke JK (1992) ALCOVE: An exemplar-based connectionist model of category learning. *Psychol. Rev.* 99: 22–44.

Larkin J, McDermott J, Simon DP, and Simon HA (1980) Expert and novice performance in solving physics problems. *Science* 208: 1335–1342.

Lassaline ME and Murphy GL (1996) Induction and category coherence. *Psychon. Bull. Rev.* 3: 95–99.

Love BC (2003) The multifaceted nature of unsupervised category learning. *Psychon. Bull. Rev.* 10: 190–197.

Love BC and Gureckis TM (2005) Modeling learning under the influence of culture. In: Ahn WK, Goldstone RL, Love BC, Markman AB, and Wolff P (eds.) *Categorization Inside and Outside the Laboratory: Essays in Honor of Douglas L Medin*, pp. 229–248. Washington, DC: American Psychological Association.

Love BC, Markman AB, and Yamauchi T (2000) Modeling classification and inference learning. *Proceedings of the Fifteenth National Conference on Artificial Intelligence* 15: 136–141.

Love BC, Medin DL, and Gureckis TM (2004) SUSTAIN: A model of human category learning. *Psychol. Rev.* 111: 309–332.

Lynch EB, Coley JD, and Medin DL (2000) Tall is typical: Central tendency, ideal dimensions and graded category structure among tree experts and novices. *Mem. Cognit.* 28: 41–50.

Malt BC, Ross BH, and Murphy GL (1995) Predicting features for members of natural categories when categorization is uncertain. *J. Exp. Psychol. Learn. Mem. Cogn.* 21: 646–661.

Markman AB and Makin VS (1998) Referential communication and category acquisition. *J. Exp. Psychol. Gen.* 127: 331–354.

Markman AB and Ross BH (2003) Category use and category learning. *Psychol. Bull.* 129: 592–613.

McRae K (2004) Semantic memory: Some insights from feature-based connectionist attractor networks. In: Ross BH (ed.) *The Psychology of Learning and Motivation: Advances in Research and Theory,* Vol. 45, pp. 41–86. San Diego: Academic Press.

Medin DL (1989) Concepts and conceptual structure. *Am. Psychol.* 44: 1469–1481.

Medin DL and Schaffer MM (1978) Context theory of classification. *Psychol. Rev.* 85: 207–238.

Medin DL, Lynch EB, Coley JD, and Atran S (1997) Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognit. Psychol.* 32: 49–96.

Medin DL, Lynch EB, and Solomon KO (2000) Are there kinds of concepts? *Annu. Rev. Psychol.* 51: 121–147.

Medin DL, Ross N, Atran S, Cox D, Coley J, Proffitt J, and Blok S (2006) Folkbiology of freshwater fish. *Cognition* 99: 237–273.

Middleton EL, Wisniewski EJ, Trindel KA, and Imai M (2004) Separating the chaff from the oats: Evidence for a conceptual distinction between count noun and mass noun aggregates. *J. Mem. Lang.* 50: 371–394.

Minda JP and Ross BH (2004) Learning categories by making predictions: An investigation of indirect category learning. *Mem. Cognit.* 32: 1355–1368.

Minda JP and Smith JD (2001) Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *J. Exp. Psychol. Learn. Mem. Cognit.* 27: 775–799.

Mufwene S (1984) The count–mass distinction and the English lexicon. In: Testen D, Mishra V, and Drogo J (eds.) *Papers from the Parasession on Lexical Semantics*, pp. 200–221. Chicago: Chicago Linguistic Society.

Murphy GL (2002) *The Big Book of Concepts*. Cambridge, MA: MIT Press.

Murphy GL (2005) The study of concepts inside and outside the laboratory: Medin versus Medin. In: Ahn WK, Goldstone RL, Love BC, Markman AB, and Wolff P (eds.) *Categorization Inside and Outside the Laboratory: Festschrift in Honor of Douglas L Medin*, pp. 197–213. Washington, DC: American Psychological Association.

Murphy GL and Allopenna PD (1994) The locus of knowledge effects in concept learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 20: 904–919.

Murphy GL and Medin DL (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92: 289–316.

Murphy GL and Ross BH (1994) Predictions from uncertain categorizations. *Cognit. Psychol.* 27: 148–193.

Murphy GL and Ross BH (2005) The two faces of typicality in category-based induction. *Cognition* 95: 175–200.

Murphy GL and Wisniewski EJ (1989) Feature correlations in conceptual representations. In: Tiberghien G (ed.) *Advances in Cognitive Science Vol. 2: Theory and Applications*, pp. 23–45. Chichester: Ellis Horwood.

Newell A and Simon HA (1972) *Human Problem Solving*. New York: Prentice-Hall.

Nosofsky RM (1986) Attention, similarity and the identification-categorization relationship. *J. Exp. Psychol. Gen.* 115: 39–57.

Nosofsky RM (1991) Relation between the rational model and the context model of categorization. *Psychol. Sci.* 2: 416–421.

Nosofsky RM, Palmeri TJ, and McKinley SC (1994) Rule-plus-exception model of classification learning. *Psychol. Rev.* 101: 53–97.

Pazzani MJ (1991) Influence of prior knowledge on concept acquisition: Experimental and computational results. *J. Exp. Psychol. Learn. Mem. Cogn.* 17: 1416–1432.

Pretz JE, Naples AJ, and Sternberg RJ (2003) Recognizing, defining, and representing problems. In: Davidson JE and Sternberg RJ (eds.) *The Psychology of Problem Solving*, pp. 3–30. Cambridge: Cambridge University Press.

Proffitt JB, Coley JD, and Medin DL (2000) Expertise and category-based induction. *J. Exp. Psychol. Learn. Mem. Cogn.* 26: 811–828.

Rehder B and Murphy GL (2003) A knowledge-resonance (KRES) model of category learning. *Psychon. Bull. Rev.* 10: 759–784.

Rosch E and Mervis CB (1975) Family resemblances: Studies in the internal structure of categories. *Cognit. Psychol.* 7: 573–605.

Ross BH (1996) Category representations and the effects of interacting with instances. *J. Exp. Psychol. Learn. Mem. Cogn.* 22: 1249–1265.

Ross BH (1997) The use of categories affects classification. *J. Mem. Lang.* 37: 240–267.

Ross BH (1999) Post-classification category use: The effects of learning to use categories after learning to classify. *J. Exp. Psychol. Learn. Mem. Cogn.* 25: 743–757.

Ross BH (2000) The effects of category use on learned categories. *Mem. Cognit.* 28: 51–63.

Ross BH and Kennedy PT (1990) Generalizing from the use of earlier examples in problem solving. *J. Exp. Psychol. Learn. Mem. Cogn.* 16: 42–55.

Ross BH and Makin VS (1999) Prototype versus exemplar models. In: Sternberg RJ (ed.) *The Nature of Cognition*, pp. 205–241. Cambridge, MA: MIT Press.

Ross BH and Murphy GL (1996) Category-based predictions: Influence of uncertainty and feature associations. *J. Exp. Psychol. Learn. Mem. Cogni.* 22: 736–753.

Ross BH, Gelman SA, and Rosengren KS (2005) Children's category-based inferences affect classification. *Brit. J. Dev. Psychol.* 23: 1–24.

Shepard RN (1962a) Analysis of proximities: Multidimensional scaling with an unknown distance function I. *Psychometrika* 27: 125–140.

Shepard RN (1962b) Analysis of proximities: Multidimensional scaling with an unknown distance function II. *Psychometrika* 27: 219–246.

Silver EA (1979) Student perceptions of relatedness among verbal problems. *J. Res. Math. Educ.* 10: 195–210.

Silver EA (1981) Recall of mathematical problem information: Solving related problems. *J. Res. Math. Educ.* 12: 54–64.

Smith JD and Minda JP (1998) Prototypes in the mist: The early epochs of category learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 24: 1411–1436.

Smith JD and Minda JP (2000) Thirty categorization results in search of a model. *J Exp. Psychology: Learn. Mem. Cognit.* 26: 3–27.

Smith JD and Minda JP (2001) Journey to the center of the category: The dissociation in amnesia between categorization and recognition. *J. Exp. Psychol. Learn. Mem. Cogn.* 27: 984–1002.

Soja NN, Carey S, and Spelke S (1991) Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition* 38: 179–211.

Solomon KO, Medin DL, and Lynch EB (1999) Concepts do more than categorize. *Trends Cogn. Sci.* 3: 99–105.

Storms G (2004) Exemplar models in the study of natural language concepts. In: Ross BH (ed.) *The Psychology of Learning and Motivation,* Vol. 45, pp. 1–39. San Diego: Academic Press.

Tomlinson MT and Love BC (2006) Learning abstract relations through analogy to concrete exemplars. In: *Proceedings of the Twenty-Eighth Annual Cognitive Science Society*, pp. 2269–2274 Hillsdale, NJ: Erlbaum.

Tversky A (1977) Features of similarity. *Psychol. Rev.* 84: 327–352.

Vigliocco G, Butterworth B, and Garrett MF (1996) Subject-verb agreement in Spanish and English: Differences in the *role* of conceptual constraints. *Cognition* 61: 261–298.

Wattenmaker WD (1992) Relational properties and memory based category construction. *J. Exp. Psychol. Learn. Mem. Cogn.* 15: 282–304.

Wattenmaker WD, Dewey GI, Murphy TD, and Medin DL (1986) Linear separability and concept learning: Context, relational properties, and concept naturalness. *Cognit. Psychol.* 18: 158–194.

Wierzbicka A (1988) *The Semantics of Grammar*. Amsterdam: John Benjamins.

Wisniewski EJ and Medin DL (1994) On the interaction of theory and data in concept learning. *Cognit. Sci.* 18: 221–281.
Wisniewski EJ, Imai M, and Casey L (1996) On the equivalence of superordinate concepts. *Cognition* 60: 269–298.
Wisniewski EJ, Lamb CA, and Middleton EL (2003) On the conceptual basis for the count and mass noun distinction. *Lang. Cogn. Process.* 18: 583–624.
Wittgenstein L (1953) *Philosophical Investigations. Anscombe GEM (trans.)*. Malden, MA: Blackwell.
Yamauchi T and Markman AB (1998) Category learning by inference and classification. *J. Mem. Lang.* 39: 124–148.
Yamauchi T and Markman AB (2000) Learning categories composed of varying instances: The effect of classification, inference, and structural alignment. *Mem. Cognit.* 28: 64–78.