

Category Dimensionality and Feature Knowledge: When More Features Are Learned as Easily as Fewer

Aaron B. Hoffman and Gregory L. Murphy
New York University

Three experiments compared the learning of lower-dimensional family resemblance categories (4 dimensions) with the learning of higher-dimensional ones (8 dimensions). Category-learning models incorporating error-driven learning, hypothesis testing, or limited capacity attention predict that additional dimensions should either increase learning difficulty or decrease learning of individual features. Contrary to these predictions, the experiments showed no slower learning of high-dimensional categories; instead, subjects learned more features from high-dimensional categories than from low-dimensional categories. This result obtained both in standard learning with feedback and in noncontingent, observational learning. These results show that rather than interfering with learning, categories with more dimensions cause individuals to learn more. The authors contrast the learning of family resemblance categories with learning in classical conditioning and probability learning paradigms, in which competition among features is well documented.

Keywords: category learning, concepts, error-driven learning, unsupervised learning

Natural concepts can be extremely rich. The amount that some people know about cars, rock music, birds, or heart disease is tremendous. All of this information can be used to classify entities into these categories or can be inferred from category membership. A bird expert can identify a bird using a few disparate properties, such as small, blue, and seen in Arizona in February (classification), and a physician can predict many symptoms of a person diagnosed with congestive heart failure (inference). Even ordinary people have vast knowledge about cats and dogs, cars, political categories, personality types, and computers. Indeed, this accumulation of knowledge is a striking accomplishment of the human conceptual system.

How is it that people acquire all of this information? There are a number of theories of concept acquisition, but, we argue, they do not give a clear answer to this question. Indeed, they suggest that learning many properties for a given concept should be extremely difficult. Furthermore, virtually all experiments on concept learning use simple categories with few stimulus dimensions. Thus, an examination of the issue of *category dimensionality*, as we call it, is overdue. In this article we first review what is known about learning concepts having few versus many properties, and then we report three experiments that contrast the learning of lower- versus higher-dimensional categories.

Aaron B. Hoffman and Gregory L. Murphy, Department of Psychology, New York University.

This work was supported by National Institute of Mental Health Grant MH41704. We thank Bob Rehder and Brian Ross for helpful comments on a draft of this article.

Correspondence concerning this article should be addressed to Aaron B. Hoffman, Department of Psychology, New York University, 8th floor, 6 Washington Place, New York, NY 10003. E-mail: aaron.hoffman@nyu.edu

Number of Dimensions in Natural and Experimental Categories

Categories studied in psychology experiments are often quite simple, in the sense of having few stimulus dimensions. Experimental categories can have as few as two (Heit, 1994; Maddox, 1995; Nosofsky, 1986) or three dimensions (Shepard, Hovland, & Jenkins, 1961; plus its many later replications). In the exemplar-prototype debate, it is common for experimental categories to have four dimensions (Medin & Schaffer, 1978; Medin & Schwanenflugel, 1981; Smith & Minda, 2000). A recent, albeit atypical study actually used only one stimulus dimension (Stewart, Brown, & Chater, 2002).

Intuitively, it seems that natural categories generally have many more dimensions than experimental categories do. But is there any way to document this difference? One can gain an idea of how many dimensions real objects have by examining feature lists. Ignoring the fact that some categories could be associated to multiple features on a dimension (e.g., that apples are green as well as red), feature lists can provide a rough estimate of the number of dimensions associated with a category. If one looks at the feature counts collected by Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976), one finds numbers such as the following (see discussion in Murphy, 2003): Sports cars have 17 features in common; sugar maples have 12; claw hammers have 8. Malt and Smith (1982) gave subjects 75 s to list properties for each of 35 category names. The mean number of properties listed by each subject for each item ranged from 6.0 to 9.6 for different birds and kinds of furniture.

Although these counts, which range from about 6 to 17 features, are already larger than the numbers of dimensions in most experimental categories, it is well known that such lists underestimate the number of features that people know of a category, because subjects tend to omit properties that seem obvious or that do not distinguish the items in the stimulus set (Tversky & Hemenway,

1984). For example, we found it very easy to list (in one sitting, about a minute) far more than the eight properties Malt and Smith's (1982) subjects listed on average for robins: flies, has wings, has a head, has two eyes, has ears, has a beak, has feathers, male has red breast, dark head, has two feet, has a tail, about 7 in. in length, migrates south in winter, eats worms and bugs, common on suburban lawns, lays eggs, egg is blue, lives in nest, nest is in trees, and has a piercing call (21 features). We could have added dozens more properties that are found in birds and animals in general—such as eats, sleeps, drinks, has a liver, has a blood, has DNA, and so forth—which apparently were not listed by subjects. It is not surprising that people do not list everything they know about robins in 75 s, especially when it is one of many items in the experiment, but a list of *everything* that one knows about robins and other common categories could clearly contain over 100 properties.

Studies of Category Dimensionality

It is natural to ask whether the results of very simple experimental categories scale up in a way that will explain how people learn categories with many more properties. Surprisingly, there is evidence in the experimental literature that it might be very difficult to learn categories with many features or dimensions (the two are usually perfectly correlated in experimental category structures).

The classic study of Shepard, Hovland, and Jenkins (1961) compared categories formed by six different rules. Although their stimuli were always constructed from three dimensions, the number of dimensions relevant to category membership varied from one to three, depending on the rule. They found a pattern in which categories based on one dimension (e.g., black vs. white shapes) were easier to learn than categories that involved two relevant dimensions (e.g., black and small or white and large), which in turn were easier than categories where all three dimensions were relevant. Shepard et al. interpreted their results as indicating that as the categories required more dimensions, learning became more difficult (e.g., pp. 4, 33). This interpretation is not entirely straightforward, however, because as the number of relevant dimensions increased, the rule relating them changed as well. For example, the unidimensional versus two-dimensional rules for one category were not, say, black versus black and large. Instead, the two-dimensional rule had two conjunctions in it: black and large or white and small. (This was necessary to retain equal category sizes across rules.) Similarly, the three-dimensional rule was not a simple conjunction of values on three dimensions, like black-large-triangle. Thus, although it is clear that Shepard et al.'s higher-dimensional rules are more difficult, it is not so clear that it is simply the addition of more dimensions that caused the difficulty. (Also, see Kruschke's, 1993, discussion of filtration vs. condensation tasks.)

The traditional concept acquisition literature of the 1960s did make a simpler comparison of rule dimensionality, generally finding that a unidimensional rule, like red things, was easier to acquire than a conjunctive rule, like red AND small things (e.g., Bruner, Goodnow, & Austin, 1956). However, most natural categories do not seem to involve such rules. Instead, they tend to follow a family resemblance structure, in which many different attributes are associated with the category, albeit not perfectly

(Rosch, 1973; Rosch & Mervis, 1975). In such a structure, attributes tend to be positively correlated with other category attributes such that they provide some redundant and some independent information. For example, imagine that a child had learned that birds are flying creatures with wings. Then, when learning that birds also have a beak and live in nests, the child is simply adding more information, which is largely redundant with respect to identifying the category. These additional dimensions do not form a more complicated rule. However, the additional information may be important in learning what birds are and what they do, even if it does not provide much higher accuracy in identifying things as birds. This is in striking contrast to Shepard et al.'s (1961) category structures, where rule complexity increased with number of relevant dimensions and individual dimensions were often not predictive of category membership. For example, for Shepard et al.'s most difficult condition (rule VI), one can classify only at chance with two dimensions but can classify perfectly when all three dimensions are fully learned.

In the present study we investigated the effect of category dimensionality on two aspects of learning family resemblance categories, addressed in the next two sections. First, what effect will increasing the number of dimensions have on the initial learning of categories, as measured by accuracy in classifying category members? Second, what effect will this variable have on learning how each dimension is related to the categories?

Initial Category Learning

Learning models differ in their predictions regarding the influence of category dimensionality on initial category learning. For instance, hypothesis-testing algorithms such as the rule-plus-exception model, RULEX (Nosofsky, Palmeri, & McKinley, 1994), propose that categorizers first test single-dimension rules and then more dimensions as needed. But the number of possible rules increases exponentially with the number of dimensions. To illustrate, there are 10 possible one- or two-dimensional rules for categories with four binary dimensions but 36 possible one- or two-dimensional rules for those with eight binary dimensions. If a category requires learning of three or more dimensions (as in our experiments), the hypothesis-testing rule would take much longer to discover this category with eight dimensions, because it would have to test and reject 26 more rules. Thus, hypothesis-testing models predict that adding dimensions should slow the learning of multidimensional categories.¹

Whereas hypothesis-testing models predict a slower learning rate for complex categories, associationist learning algorithms (e.g., Gluck & Bower, 1988; Kruschke, 1992) predict the opposite. Networks can accommodate additional dimensions with more input nodes; this yields a larger total amount of activation delivered to the category output, resulting in a higher percentage of correct responses. Thus, unconstrained networks can use the additional informative dimensions in a large family resemblance category structure to yield faster learning.

¹ RULEX combines rule learning with exemplar learning, and it has parameters that vary when the learner shifts from the former to the latter. Thus, exact predictions for RULEX depend on the values of its parameters. We are discussing here hypothesis testing alone, as in RULEX's first component.

Network models can, however, assume an attentional cost associated with adding dimensions. RASHNL (Kruschke & Johansen, 1999), for example, normalizes the weights connecting input nodes to the rest of the network so that the total influence of the input is limited in capacity (i.e., attentional weights must sum to 1.0). RASHNL predicts that increasing the number of dimensions (complexity) will spread attention more broadly across the inputs. Spreading attention could reduce and even counteract the positive influence of adding informative dimensions. If the attentional capacity is large, adding dimensions will likely facilitate category learning. If attentional capacity is small, adding dimensions may produce no change in the learning rate, as increased input activation is offset by lower attentional weights. Thus, depending on the exact details of a model's learning algorithm and its attentional assumptions (and their interaction), a model might predict a positive or negative effect of the number of dimensions, or it might even predict that attentional limits and the benefit of increased activation (from more dimensions) cancel each other out, resulting in no change in learning rate.

More broadly, then, it is clear that hypothesis-testing and associationist accounts make different predictions regarding category complexity. Whereas hypothesis testing predicts a decrement in the learning rate with added dimensions, associationist accounts can predict no effect or facilitation. However, within associationist accounts, predictions can vary depending on their assumptions regarding attentional capacities.

Feature Learning

A second important issue is how category complexity affects the learning of the category's features. Here the predictions are fairly straightforward and are based on how increasing the number of dimensions should cause interference between features. Indeed, such *cue competition effects* are found across learning situations and animal species, suggesting that more dimensions will reduce the amount learned about any particular dimension (see Kruschke & Johansen, 1999). One paradigm demonstrating cue competition is nonmetric multiple-cue probability learning, in which properties (and items) are associated with category labels only probabilistically, so that perfect performance is never attainable. This is similar to family resemblance categories in that features are predictive of category membership but not perfectly so. Researchers have found in this task that any particular dimension will be used less in classification when additional irrelevant dimensions are provided (Edgell et al., 1996). This competition effect increases with the salience and cue validity of the additional dimensions (Edgell, Bright, Ng, Noonan & Ford, 1992; Edgell et al., 1996, Experiment 6). Thus, although adding more dimensions to a category might not harm exemplar classification, cue competition should reduce how much a dimension is used and, therefore, the strength of the association between that dimension and the category.

Cue competition may be traced to basic processes in associative learning, as reflected in the phenomena of blocking and overshadowing in classical conditioning tasks (Pearce & Bouton, 2001). When one cue or set of cues already predicts a given outcome, learning of new cues is reduced. Error-driven learning algorithms such as the delta rule (Widrow & Hoff, 1960) common in connectionist learning models also have this property (see Gluck &

Bower, 1988; Kruschke, 1993): Because learning is propagated through error, if one cue or set of cues successfully predicts the outcome, there is little error, and therefore the association weights of other cues do not change with further experience (see Rehder & Murphy, 2003, for an example).

Another mechanism by which increased dimensions might negatively influence feature learning is attentional limits, which have been the focus of much theorizing in category learning (Kruschke & Johansen, 1999; Nosofsky, 1986; Rehder & Hoffman, 2005a). As explained in the previous section, attention can limit the maximum input that the system can process at once. This is implemented, for example, by normalizing the weights of a network's input nodes. Under the assumption that attention is limited, adding more dimensions must spread attention (i.e., weights) more thinly, thereby decreasing the learning of any specific dimension. Thus, even if learning of the category as a whole is not slowed, learning of individual dimensions ought to diminish when more dimensions are added.

The Present Research

In short, there are many reasons to think that categories with more dimensions may be disadvantaged, in two respects. First, in the acquisition of the category, as defined by correct classification of learning exemplars, learning may be slowed (although this depends on the learning model, as discussed above). Second, the different properties of more complex categories may interfere with one another and compete for attention, thereby causing less learning on average for each dimension.

The puzzle, then, is why it seems that people find it quite easy to learn natural categories, including their dozens of features. Indeed, the accepted wisdom in the conceptual development literature is that children can learn much about a category after viewing one or two exemplars, and it is apparent from word learning that young children are acquiring many lexical concepts every day (see, e.g., Bloom, 2000; Carey, 1978; Murphy, 2001). In contrast, adult subjects in psychology experiments often take many blocks to learn artificial concepts with significantly fewer dimensions. There are several likely determinants of the slowness of category learning in experiments, but it is nonetheless striking that the richness of natural categories does not seem to make them difficult to acquire.

We have so far characterized the issue of category complexity as one that has been virtually ignored. One notable exception is Minda and Smith's (2001) Experiment 4, which compared four- and eight-dimensional categories in tests of whether learners' categorization was better characterized by a prototype or an exemplar model. Contrary to a number of predictions reviewed above, they found that accuracy during learning was similar for the two category structures: Average accuracy was 79% and 76% correct for the four-dimensional and eight-dimensional structures, respectively. We do not know, however, whether subjects in the high-dimensional condition learned less about individual features than those in the low-dimensional condition, because Minda and Smith did not perform tests on individual feature knowledge.

The goal of the present study, therefore, was to investigate the variable of the number of dimensions in family resemblance category learning. The results are important both from a descriptive perspective, as this is an important variable that has not been sufficiently studied, and from a theoretical perspective, as the

results will speak to the attentional and learning mechanisms that make predictions about category complexity.

Experiment 1

We have described the comparison we are making as one of “adding dimensions” to a category, which is useful shorthand for our manipulation. However, when one adds dimensions, one is often changing other things about the category. Some of these other things are inherent differences between rich and simple categories, and they should not be “controlled away.” Other things may be somewhat independent, and so across three experiments we tested two category structures that differed in exactly how dimensionality varied. We attempt to specify in what respects the simple and rich categories differ in each experiment. However, we do not believe there is one single correct way to vary category complexity—rather, we suspect that each experiment provides a different kind of comparison that likely exists somewhere in the world. Each way of adding dimensions provides part of the answer to the question of how simple and complex categories differ.

In the present experiment, we started with a simple four-dimensional family resemblance category (which we call the *4-d structure*, following the popular one-away design (from Medin, Wattenmaker, & Hampson, 1987, and many subsequent experiments), as shown on either side of Table 1. Each column represents a different stimulus dimension, and the 1s and 0s refer to values on the dimensions. The stimuli in these experiments were pictures of bugs (referred to as Mobbles and Streaths), and the dimensions refer to parts of the bugs, such as their wings, eyes, feet, and various markings on their bodies. So, if Dimension 1 is the bug’s eyes, then a 0 might indicate two eyes, and a 1 four eyes. As Table 1 shows, each dimension has a value that is generally associated with a category. Mobbles are associated with 0s, and Streaths with

1s. That is, Mobbles usually (but not always) have two eyes, and Streaths usually have four eyes.

To increase the number of category dimensions, in the eight-dimensional (*8-d*) condition of Experiment 1, we simply doubled the 4-d structure (i.e., using all eight dimensions of Table 1). That is, the first and last four dimensions of the 8-d condition are structurally identical. This structure equates the cue validity of each dimension so that each dimension is equally predictive of category membership. Furthermore, loosely speaking, the new dimensions did not have to be learned. That is, if subjects learned Dimensions 1–4, their accuracy in categorization would not be increased by learning Dimensions 5–8.

To measure category coherence for the two conditions, we calculated the ratio of within-category similarity (number of shared features between category exemplars) to between-category similarity (see Minda & Smith, 2001). One effect of doubling the 4-d structure to construct the 8-d structure was that the number of exception features within each category doubled (increasing between-category similarity), but so did the number of shared features within categories. As a result, category coherence was equal, at 2.13, for the two conditions.

However, adding four dimensions in the 8-d condition might affect how people learn the categories. If subjects were attempting to form a rule to learn the categories, the addition of the new dimensions is not completely neutral. To understand why, first consider the 4-d structure. Subjects cannot learn these categories by learning one or two dimensions, because of the *exception features* in the design (1s in Mobbles and 0s in Streaths). If one has learned that Mobbles have two eyes and horizontal stripes, one will be baffled by one item that has four eyes and horizontal stripes and another that has two eyes and vertical stripes. Therefore, one must learn a third dimension (e.g., tail shape) in order to break ties of this sort.

The same is generally true of the rich category structure with eight dimensions, except that now it makes a difference which dimensions one happens to learn first. For example, if one learned Dimensions 1, 6, and 8, one could accurately identify all of the category members. But if one instead learned Dimensions 1, 5, and 8, one could not. The reason is that when Dimension 1 has an exception feature, so does Dimension 5, and thus one would need to learn five total dimensions in order to outweigh these two exception features and choose the correct category. In short, some combinations of three dimensions will result in correct classification of all items, but others will not. We calculated the probabilities that subjects would require three or five dimensions to correctly classify the items in the 8-d condition (assuming random selection of dimensions) and discovered that they are .57 and .43, respectively. (That is, random selection of three dimensions would result in two redundant dimensions .43 of the time.) Therefore, if subjects are attempting to learn rules to classify the items (which is not to be taken for granted), they would need to learn 3 dimensions in the 4-d case but 3.9 dimensions (.57 × 3 + .43 × 5) in the 8-d case, on average. Thus, one might well expect the 8-d categories to take longer to learn.

In Experiment 3, we changed the category structure to separate the effect of number of dimensions from this aspect of the category-learning rule. However, it is not clear that the greater difficulty of forming rules for richer categories is not present in natural categories. That is, similar considerations might influence

Table 1
The 4-d and 8-d Category Structures, Experiments 1 and 2

Stimulus	8-d							
	4-d ₁				4-d ₂			
	D1	D2	D3	D4	D5	D6	D7	D8
Mobbles								
1	0	0	0	1	0	0	0	1
2	0	0	1	0	0	0	1	0
3	0	1	0	0	0	1	0	0
4	1	0	0	0	1	0	0	0
5	0	0	0	0	0	0	0	0
Streaths								
1	1	1	1	0	1	1	1	0
2	1	1	0	1	1	1	0	1
3	1	0	1	1	1	0	1	1
4	0	1	1	1	0	1	1	1
5	1	1	1	1	1	1	1	1

Note. The 4-d conditions contained either Dimensions 1–4 or Dimensions 5–8. The 8-d conditions included all eight dimensions. The 1s and 0s refer to values on a dimension (D).

the learning of complex concepts such as birds or cars as opposed to balls or stones.

In Experiment 1 we compared the learning of 4-d and 8-d concepts in a standard category-learning paradigm. After the learning phase, we tested subjects' knowledge of individual features, to discover whether the presence of many stimulus dimensions had interfered with the learning of any one of them. Following most previous experiments on category learning, we used a learning criterion to determine when the learning phase would stop and the test phase begin. This allowed us to track differences in learning, but it caused some difficulty in interpreting the later test results, because different subjects had different amounts of exposure to category members (although we addressed this possibility in our analysis). Later experiments used a fixed learning stage, which allowed us to avoid this concern.

A final technical matter concerns just how to compare what has been learned in the two conditions. A typical measure might be proportion correct in the feature test. However, that measure has the problem that it is dependent on the arbitrary factor of how many dimensions we decided to put into each category. If we had made complex categories with 12 or 20 or 100 dimensions, then we would likely be requiring subjects to learn much more than could reasonably be learned during an experimental session. Furthermore, if subjects actually did learn more dimensions in such large-dimensional categories, this would be masked by the fact that proportion correct divides by the total number of dimensions. (For example, if subjects learned 2 out of 4 dimensions in the simple condition and 8 out of 20 dimensions in the complex condition, the proportion correct would be lower in the latter condition, even though subjects learned four times as many dimensions.) Thus, although we report the proportions correct in each condition, the more relevant statistic seems to be the number of dimensions learned. That is, does adding more dimensions actually reduce the number learned, through interference?

Method

Subjects. Twenty-four New York University students participated for pay. Half were randomly assigned to the 4-d and half to the 8-d condition.

Materials. The stimuli were 10 drawings of bugs labeled "Mobbles" and "Streaths." The category structures were composed of stimuli with four or eight binary dimensions. In the 4-d condition, subjects learned the category structure shown in Table 1 under 4-d₁ or 4-d₂. As represented in Table 1, the Mobble category had 0 as the most common value on each dimension, whereas Streaths had 1 as the most common value on each dimension.

Table 1 also shows the category structure for the 8-d condition, consisting of Dimensions 1 through 8. Thus, there were two exception features per item in all but the two prototypes. The eight dimensions of the bugs were as follows: number of eyes (two or four), antennae (straight or curly), wing (rounded or angular), upper body stripes (vertical or horizontal), foot pad (orange circle or purple triangle), lower body mark (white or blue), stinger (single or double), and front feet (claw or furry).

As Table 1 indicates, we created two versions of the 4-d condition, one using Dimensions 1–4 and the other using Dimensions 5–8, so that the 4-d and 8-d conditions had the same average feature salience. We roughly balanced the two 4-d categories by obtaining informal salience judgments and then distributing the eight dimensions to minimize salience differences between the two sets of dimensions. There were no reliable differences in any of the reported experiments between the two 4-d categories. Also, we have since replicated our results in two other experiments (not reported

here) with different stimulus dimensions,² and so stimulus-specific variables do not seem to be a factor. Feature-category label assignment was the same across subjects.

Procedure. We presented items randomly in blocks of 10. A trial consisted of an image of a single bug appearing in the center of a computer monitor, to which subjects responded by pressing *z* or */* on the keyboard. After the response, the correct category label appeared above the bug, and the word *CORRECT* or *INCORRECT* appeared below it, for 4 s. Subjects continued classifying bugs until they responded perfectly for a block or completed 30 blocks. At the end of each block, subjects were informed how many items they had correctly classified.

After the learning phase ended, the transfer phase began. Subjects viewed one of the learning items or a single-feature item (as described below) and classified it as quickly as they could. No feedback was given in this phase. Training and single-feature items were presented twice, resulting in 36 test trials for the 4-d condition and 52 for the 8-d condition.

Single-feature transfer tests. The purpose of the single-feature transfer tests was to assess subjects' category knowledge of each individual dimension. To this end, we constructed test items that displayed only a single feature. The question of interest is how many dimensions each subject learned across conditions. However, simply by guessing, subjects in the 8-d condition would appear to have learned more dimensions, because they had more dimensions on which to guess. Therefore, we applied the following guessing correction to estimate the number of dimensions each subject learned: $D_{\text{learned}} = D_{\text{total}}(P_{\text{correct}} - P_{\text{incorrect}})$, where D_{total} is the total number of dimensions in the learned category and P_{correct} and $P_{\text{incorrect}}$ are the proportions of correct and incorrect responses to single-feature items, respectively. If half of a subject's guesses are correct and the other half incorrect, then subtracting the incorrect guesses from the correct responses will leave only the known correct answers (on average). If subjects solely guessed, they would have as many incorrect as correct answers and would get a score of 0. If they learned all of the features correctly, they would receive scores of 4 and 8 in the 4-d and 8-d conditions, respectively. In an alternative measure of individual dimension knowledge, subjects were recognized as having knowledge of a dimension only if they responded correctly on all four tests of that dimension during transfer. This is a strict criterion, because it does not allow response error, but it also makes no assumptions regarding guessing.

Results

The dependent variables were errors during learning and in the transfer phase, and reaction times (RTs). Error data are the most common measure in the field, and, indeed, RTs were generally quite variable. Moreover, the comparison of 4-d and 8-d stimuli is not interpretable for RTs, because the stimuli differed greatly in perceptual complexity, and so any differences cannot be clearly attributed to differences in category knowledge. As a result, we do not report the RTs for learning and whole-item transfer tests but do so for the single-feature tests, where the items are identical in the two conditions.

Learning. We first compared how well subjects learned the category structures across conditions. Almost all subjects obtained perfect performance by the 30th block: There was only one non-learner in the 4-d condition and two in the 8-d condition. We gave those subjects who completed 30 blocks without reaching perfect performance a score of 31. The average number of blocks to criterion was numerically smaller in the 4-d condition ($M = 11.0$,

² One experiment used similar bug stimuli but with entirely different features, and the second experiment used textual features having to do with categories of vehicles.

$SD = 10.3$) than in the 8-d condition ($M = 15.2$, $SD = 8.6$), but this difference was not statistically reliable, $t(22) = 1.1$, $p > .25$. Of course, not finding reliable differences does not mean the conditions were equally easy. In fact, a power analysis indicated little chance (20%) of finding statistical reliability with alpha at .05 had the difference of four blocks been real. (To obtain 80% power would have required an additional 65 subjects per condition.) Thus, we cannot make too much of the absence of a learning difference; the next two experiments addressed this issue further. However, the apparent difference between conditions were to some degree caused by nonlearners, who received scores of 31 blocks, increasing variance considerably. If they are excluded, variability indeed declines, but the mean difference also declines to 2.8 ($M_s = 9.2$ and 12.0 blocks, $SD_s = 8.6$ and 4.8 , for the 4-d and 8-d conditions, respectively), which is still not reliable ($t < 1.0$).

Whole-item transfer. To assess what was learned about the categories, we compared how well subjects performed in classifying the whole items without feedback. It appeared that one subject who reached the learning criterion in the 8-d condition reversed responses in the transfer phase and therefore performed far below chance. We reversed his responses in the following analyses. The average proportions correct were approximately equal, .87 ($SD = .11$) and .89 ($SD = .12$) for the 4-d and 8-d conditions, respectively ($t < 1.0$). Thus, after the learning phase, subjects in the two conditions demonstrated equal knowledge of the training items.

Single-feature transfer. Subjects yielded similar proportions of correct responses for the single-feature items in the 4-d and 8-d conditions ($M_s = .82$ and $.80$, $SD_s = .14$ and $.10$; $t < 1.0$). Using the guessing correction described above, we estimated the number of dimensions learned in the 8-d condition and found it higher on average ($M = 4.8$, $SD = 1.7$) than the number of dimensions learned in the 4-d condition ($M = 2.6$, $SD = 1.1$), $t(22) = 3.7$, $p < .01$. Indeed, 6 of 12 subjects in the 8-d condition learned five or more dimensions (2 subjects learned seven dimensions and 1 learned six dimensions), whereas only 3 of 12 subjects in the 4-d condition learned all four dimensions. Moreover, although the averages were reduced, the pattern of results was similar with the strict perfect-performance criterion: Subjects in the 8-d condition identified 4.3 ($SD = 2.1$) dimensions perfectly on average, whereas those in the 4-d condition identified 2.5 ($SD = 1.2$) dimensions perfectly, $t(22) = 2.6$, $p < .05$. With the strict criterion there were still 6 subjects in the 8-d condition who learned five or more dimensions (2 learned seven and 1 learned six) but only 3 in the 4-d condition who learned all four dimensions. Thus, subjects demonstrating knowledge of additional dimensions did so with a high level of consistency.

The superior dimension knowledge in the 8-d condition compared with the 4-d condition was probably not due to speed-accuracy trade-offs because the average RT to classify single features in the 4-d condition ($M = 1,978$, $SD = 1,208$) did not differ statistically from that of the 8-d condition ($M = 2,232$, $SD = 617$; $t < 1.0$). There was, however, an outlying subject in the 4-d condition who may have increased the average RT for this condition. We first log transformed RTs to reduce the effects of outliers. The transformed data, however, still failed to produce a reliable difference between the two conditions. Only after removing the outlier was the 4-d condition statistically faster than the 8-d condition, $t(21) = 2.6$, $p < .05$. Thus, the RT results here are not entirely clear. This potential difference did not replicate in later experiments. Also, the correlation between RT and errors did not

approach significance, $r(22) = -.20$, $p > .10$, further arguing against a speed-accuracy trade-off.

Because learners in the 8-d condition took four blocks more on average to reach the learning criterion, an alternative (or additional) explanation for the results could be differences in exposure to the exemplars (although learning differences were not reliable). We addressed this potential confound by adjusting statistically for differences in number of blocks in a regression analysis—effectively analyzing dimension knowledge while statistically equating blocks to criterion. If greater exposure were the reason for the 8-d condition's advantage in feature learning, then blocks to criterion should predict number of dimensions learned. By regressing the number of dimensions learned on category complexity and, simultaneously, number of blocks to criterion, we contrasted the partial predictive power of each variable. Blocks to criterion was mean-centered to allow interpretability of its interaction with condition (see Judd & McClelland, 1989, p. 258). We estimated the following regression equation, where *dimensionality* refers to the 4-d versus 8-d comparison: dimensions learned = $2.5 - .02 * (\text{blocks to criterion}) + 2.5 * (\text{dimensionality}) - .1 * (\text{blocks} \times \text{dimensionality})$. The effect of blocks to criterion was negative ($-.02$), that is, in the opposite direction than expected if increased exposure had caused the 8-d advantage in feature knowledge. Moreover, it was not reliable ($t < 1.0$), indicating that the additional blocks to criterion for the 8-d condition did not by itself improve learning. Consistent with our previous analysis, the estimated difference in number of dimensions learned between the 8-d and 4-d conditions was 2.5, $t(20) = 4.7$, $p < .01$. The interaction was negative and marginally reliable, $t(20) = 1.8$, $p < .10$. For the 4-d condition, blocks to criterion had no effect on dimensions learned, but for the 8-d condition, every additional block resulted in a decrease in dimensions learned by .10. Again, this interaction is opposite to the prediction of the alternative explanation, probably reflecting the fact that poor learners took longer to reach criterion and also performed poorly on the test.

Thus, when we statistically adjusted for number of blocks to criterion, the average number of dimensions learned was 5.0 in the 8-d condition and 2.5 in the 4-d condition. The estimated number of dimensions learned for the 8-d condition *increased* from our previous estimate, because the effect of blocks (for the 8-d condition only) was in the opposite to expected direction. Inspection of the data indicated that the negative effect of blocks to criterion may have been carried by a single nonlearner in the 8-d condition. Indeed, the interaction disappeared with his removal. Removing all nonlearners resulted in an estimated 5.2 and 2.6 dimensions learned in the 8-d and 4-d conditions, respectively.

Discussion

The learning and whole-item transfer results were contrary to the idea that people engage (solely) in hypothesis testing. That is, even with a 2-to-1 ratio in number of dimensions between the 8-d and 4-d conditions, we found no reliable differences in learning difficulty or transfer accuracy. It seems, then, that these results are consistent with associationist models (with attention capacity limits), which can predict no effect of added dimensions on learning difficulty. However, as discussed above, we modified the learning procedure in later experiments, because the null effect of learning differences here is difficult to interpret.

More critical is the finding that subjects in the 8-d condition learned more dimensions than those in the 4-d condition. This result is contrary to error-driven learning and limited capacity attention, suggesting that subjects in the 8-d condition used other learning strategies or applied attention differently than those in the 4-d condition. We address attentional issues and multiple learning strategies in the General Discussion.

We also considered the possibility that the four more blocks (on average) required by subjects in the 8-d condition to reach the learning criterion provided more exposures to learn additional dimensions. To some degree this difference was driven by nonlearners, who received a score of 31 blocks (as compared with learners, who averaged about 11 blocks). The regression analysis revealed a negative and unreliable relationship between blocks to criterion and dimension knowledge. Nonetheless, to rule out any possible effect of amount of exposure during learning, we equated the number of learning blocks across conditions in Experiments 2 and 3.

Experiment 2

Equating the number of learning blocks allows us to compare the results from the transfer tests in a straightforward manner by preventing any confounding of amount of exposure and category dimensionality. Additionally, Experiment 2 tested whether subjects will once again learn more features in the 8-d than in the 4-d condition with observational learning, in which subjects simply observe category exemplars with their labels.

In the traditional category-learning paradigm, subjects may try to correctly classify items by learning which features predict which category label (e.g., if rounded wing then Mobbler). Explicit hypothesis testing of one or more dimensions is often an effective strategy for this task. In particular, learners may attempt to test a single dimension, followed by an additional dimension, and so on, until the rule is successful. At that point, learning stops (Nosofsky et al., 1994). In observational learning, on the other hand, the category name is provided initially, and the learner is therefore less likely to actively test a rule. Rather, the subject may simply try to associate the presented features or exemplar to the category label. As a result, observational learners may try to process category items as a whole instead of analytically identifying the minimal dimensions necessary for correct classification. If this is the case, then observational learners could learn more dimensions; in particular, the 8-d condition's advantage in number of dimensions learned (found in Experiment 1) might be larger for observational than for feedback learning. (See related manipulations in Markman & Ross, 2003, and Ashby, Maddox, & Bohil, 2002.)

Method

Subjects. Forty-eight New York University students participated for pay. One quarter of the subjects were randomly assigned to each of the four conditions.

Procedure. The same category structures (see Table 1) and bug stimuli from Experiment 1 were used here. The learning and transfer phases were similar to those of Experiment 1 except for two crucial differences. The first difference was that half of the subjects were assigned to an *observational* learning condition and half to the *standard* feedback learning that subjects underwent in Experiment 1. For the observational learning condition, we effectively reversed the standard learning procedure: Subjects first viewed the category label together with the bug for 4 s. The label then

disappeared, and the subjects could study the bug for an unlimited amount of time. They pressed the key corresponding to the correct category label for the stimulus to go to the next trial. The standard feedback learning condition was identical to that of Experiment 1, in which subjects first classified the bug and then received 4 s of feedback. Thus, both groups pressed a button and received category information for 4 s, but subjects in the standard condition had to classify the exemplar prior to receiving the label. The experiment was a 2 (learning condition) \times 2 (dimensionality) between-subjects design.

The second procedural difference between Experiment 1 and the current experiment was that we removed the perfect-block criterion from the learning phase. Instead, all subjects received eight learning blocks. Eight blocks was the median for subjects to reach the learning criterion in Experiment 1. Thus, by Block 8 we expected most subjects to have approached or achieved perfect performance.

Results

Learning. We first asked whether the 4-d condition's expected learning advantage would obtain in the current experiment, for the standard feedback condition only (there is no measure of accuracy for the observation learning subjects). Overall, 5 subjects in the 4-d condition and 5 in the 8-d condition achieved an errorless block at some point during learning (the criterion from Experiment 1), suggesting no category-learning advantage for the 4-d condition. The conditions showed similar and substantial levels of learning, with mean proportions correct of .72 and .71 for the 4-d and 8-d conditions over learning. Indeed, a category structure by block mixed analysis of variance revealed an effect of block, $F(7, 154) = 11.1$, $MSE = 0.02$, $p < .01$, but no effect of category complexity or its interaction with block ($F_s < 1$), indicating that classification performance was similar across the two conditions. In short, consistent with Experiment 1's results, the various measures indicate that the 4-d and 8-d category structures were approximately equally difficult to learn.

Whole-item transfer. Table 2 summarizes the results for the whole-item transfer tests. Subjects in the standard feedback condition performed 6% better than those in the observational condition on average, $F(1, 44) = 2.9$, $MSE = 0.02$, $p < .10$, probably because this test was exactly the same as the learning procedure for that condition. However, there was no difference in performance between the 4-d and 8-d conditions, nor were there any reliable interactions ($F_s < 1$). This is further evidence that category learning was equivalent in the two category structures.

Single-feature transfer. Since number of learning blocks was held constant, all differences in transfer can be attributed to experimental manipulations, rather than to differences in exposure. Table 2 summarizes the accuracy results for the single-feature items (right columns). Overall, the 8-d condition had marginally lower accuracy (a 7.4% difference), $F(1, 44) = 2.8$, $MSE = 0.02$, $p < .10$, and a similar difference (11%) was found within the feedback condition, $t(22) = 1.7$, $p < .10$. However, the 8-d condition was also marginally faster, by about 500 ms, $F(1, 44) = 3.5$, $MSE = 966,713$, $p < .10$. Thus, the 8-d subjects' marginally worse accuracy may have been caused in part by their marginally faster responding. Also, there was one unusual subject in the 8-d condition who learned by the second block yet subsequently responded close to chance for the remainder of the experiment.

As discussed earlier, the number-of-dimensions-learned measure seems a fairer test of the two conditions, because it does not divide by the arbitrary factor of how many dimensions we pre-

Table 2
Whole-Item and Single-Feature Transfer Results for Standard Feedback and Observational Learning Conditions, Experiment 2

Category structure	Whole items		Single-feature items	
	Proportion correct	Proportion correct	RT	Dimensions learned
Feedback learning				
4-d	.85	.84	2,315	2.71
8-d	.82	.73	1,986	3.71
Observational learning				
4-d	.77	.82	2,462	2.54
8-d	.78	.78	1,725	4.42

Note. RT = reaction time (in ms).

sented. The 8-d condition learned more dimensions (4.1) than the 4-d condition (2.6), $F(1, 44) = 7.4$, $MSE = 3.3$, $p < .01$. Moreover, the 8-d structure advantage for single-feature knowledge held regardless of the learning condition (interaction $F < 1$).

Thus, although the learning rates, the number of subjects who reached perfect performance, and the whole-item transfer performance were very similar across the two conditions, subjects in the 8-d condition learned more dimensions than those in the 4-d condition. Indeed, 10 of 24 subjects in the 8-d condition learned five or more dimensions (2 subjects learned all eight, and 2 learned seven), whereas just 8 subjects learned all four dimensions in the 4-d condition. With the perfect-performance criterion, subjects in the 8-d condition learned 3.9 dimensions ($SD = 2.2$), and those in the 4-d condition learned 2.5 dimensions ($SD = 1.4$), $t(46) = 2.6$, $p < .05$. With this criterion, 10 subjects learned five or more dimensions in the 8-d condition (5 learned seven and 1 learned six), whereas only 8 learned all four dimensions in the 4-d condition.

Discussion

In Experiments 1 and 2, subjects in the 8-d condition demonstrated greater dimension knowledge than did those in the 4-d condition. The 8-d advantage in feature learning cannot be attributed to additional learning blocks, because their number was held constant in Experiment 2. Moreover, although subjects presumably had to spread their attention (and error) over a greater number of dimensions, they learned more dimensions. Subjects learning the 8-d categories learned more dimensions, on average, even by our strict perfect-response criterion for dimension knowledge. Thus, limited-attention association accounts were once again contradicted by the single-feature test results.

We also hypothesized that analytic learning strategies could cause learners to focus on the minimal number of dimensions required for classification. To test this hypothesis we compared an observational learning condition (intended to reduce analytic learning strategies) with the standard learning condition. However, the greater feature learning for 8-d categories was not reliably larger in the observational learning condition.

Observation learning nevertheless provided an opportunity to test the 8-d advantage for feature learning in another learning condition. This was crucial because, as this advantage was a newly discovered phenomenon, we did not know the extent to which it applied to the wide variety of learning conditions in natural con-

cept acquisition. Thus, finding that subjects in the higher-dimensional, observational learning condition learned more dimensions supported its generalizability.

Although the feature learning advantage for the 8-d condition held across two different learning conditions, we have so far tested the effect of adding dimensions in just one category structure. Thus, we do not yet know whether this finding is a result of some aspect of that particular structure. As we noted, rule learning would require more dimensions on average in the 8-d than in the 4-d condition (3.9 vs. 3.0 dimensions). Therefore, an alternative explanation is that increasing the difficulty of the category rule caused subjects to have learned more dimensions. Indeed, only a third of the 4-d condition subjects (across experiments) learned all four dimensions, which is consistent with a three-dimensional classification rule. Thus, perhaps our result simply reflects the fact that subjects needed to learn more dimensions in the 8-d case and therefore they did so. This argument is not quite sound, however, because it should be more difficult for learners to acquire more dimensions, and yet we found that the 8-d group learned more even when the learning blocks were fixed. That is, the fact that the 8-d group needed to learn more dimensions in order to be perfect should have made it harder for them to learn the categories, which did not occur. Furthermore, the fact that more dimensions ought to be learned does not entail that more dimensions *will* be learned, when exposure to exemplars is held constant. Nonetheless, the differences in the required number of dimensions may be somehow responsible for the results so far.

Another issue with the present design is that pairs of dimensions in the 8-d condition are redundant. Therefore, one might question whether the 8-d structure really has more dimensions than the 4-d structure. If two dimensions are completely correlated, there is nothing to prevent the learner from effectively combining those two dimensions into a single dimension (though of course he or she would have to notice this correlation).

Experiment 3 used a different category structure that addressed both of these possible concerns about the structure used in Experiments 1 and 2.

Experiment 3

Experiment 3 addressed the disparity in the minimum number of dimensions required between the 4-d and 8-d conditions by replacing the structure shown in Table 1 with an eight-dimensional

“one-away” structure, shown in Table 3. Because there are no redundant dimensions, the minimum number of dimensions required for perfect classification is three for both conditions. Notice also that the 8-d structure now has eight items per category rather than five (the one-away design requires the same number of exemplars as stimulus dimensions). To equate the number of instances per learning block across category structures, we added three prototypes to the 4-d condition. The resulting 4-d category structure comprises eight items, of which half are the category’s prototype. This design also equates the cue validity of each dimension across category structures (.88), and the category coherence scores for the two structures (calculated as in Experiment 1) were both 3.57. As additional prototypes and duplicate items should favor the 4-d condition, the design is somewhat biased against the 8-d condition. This is thus a conservative test of our earlier results that subjects learn more dimensions in the 8-d condition.

The eight nonredundant dimensions served to test the two alternative hypotheses for the observed differences in feature learning. One hypothesis was that the greater number of dimensions required to learn the 8-d structure in some cases (up to five dimensions, depending on which ones are selected) led to the observed difference. As both structures can now be learned with three dimensions, this difference is no longer present. The second hypothesis was the possibility of combining the redundant dimensions into a single dimension. Because there are no redundant dimensions in either condition in Experiment 3, this difference also cannot explain the results. Indeed, this design is conceptually quite close to the classical conditioning blocking paradigm described in

the introduction. That is, once subjects have learned three dimensions, they will perform perfectly, and further learning cannot improve their performance in either condition. Therefore, if the analogy to classical conditioning is apt, or if error-driven learning explains which features are acquired, subjects should learn about three dimensions in both conditions.

Method

Subjects. Twenty-four New York University students participated for pay. Half were randomly assigned to the 4-d and half to the 8-d condition.

Materials. The stimuli were again depicted bugs. The category structures were composed of four (4-d condition) or eight (8-d condition) binary dimensional stimuli, as shown in Table 3. Thus, the expected minimum number of dimensions (three) was equivalent across the conditions. For the 4-d conditions, the structure was identical to that used in Experiments 1 and 2 with the addition of three prototype items (four prototype items in total).

Procedure. The learning and transfer phases were the same as in Experiment 2. However, given the new category structures, there was now a total of 16 items in both conditions. Training and single-feature items were presented twice, resulting in 48 test trials for the 4-d condition and 64 for the 8-d condition.

Results

Learning. Eight subjects in the 4-d condition and 11 in the 8-d condition reached perfect performance at some point during learning. Thus, after we equated the minimum dimensions required for learning, the results suggest a category-learning advantage for the 8-d condition. Figure 1 shows proportion correct over blocks for each condition. The 8-d condition yielded a higher proportion correct across learning blocks ($M_s = .86$ and $.92$ for the 4-d and 8-d conditions, respectively). The main effects of block, $F(7, 154) = 17.4$, $MSE = .005$, $p < .01$, and condition, $F(1, 22) = 11.5$, $MSE = .01$, $p < .01$, were reliable, but there was no interaction ($F < 1$).

Whole-item transfer. One subject who reached the learning criterion in the 4-d condition apparently reversed responses and therefore performed far below chance during transfer. In the following analyses we reversed her responses. The average proportions correct were high and equal ($M = .95$, $SD = .05$, and $M = .96$, $SD = .07$, for the 4-d and 8-d conditions, respectively; $t < 1.0$).

Single-feature transfer. Proportion correct in the 8-d condition ($M = .78$, $SD = .13$) was nonsignificantly less than that of the 4-d condition ($M = .86$, $SD = .12$), $t(22) = 1.6$, $p > .10$. However, the estimated number of dimensions learned was 4.5 ($SD = 2.1$) in the 8-d condition and only 2.9 ($SD = 1.0$) in the 4-d condition, $t(22) = 2.3$, $p < .05$. Moreover, 8 subjects in the 8-d condition learned more dimensions than the required three (3 subjects learned seven, 2 learned five, and 3 learned four), whereas only 3 subjects in the 4-d condition learned all four dimensions. Similarly, with the perfect-performance criterion, the number of dimensions learned was 4.3 ($SD = 2.0$) in the 8-d condition and 2.8 ($SD = 1.0$) in the 4-d condition, $t(22) = 2.5$, $p < .05$. RTs did not differ across conditions ($t < 1.0$). Thus, both criteria suggested that subjects in the 8-d condition learned more dimensions than those in the 4-d condition, even though the logical requirements for distinguishing the categories were identical.

Table 3
The 4-d and 8-d Category Structures, Experiment 3

Stimulus	8-d							
	4-d ₁				4-d ₂			
	D1	D2	D3	D4	D5	D6	D7	D8
Mobbles								
1	0	0	0	0	0	0	0	1
2	0	0	0	0	0	0	1	0
3	0	0	0	0	0	1	0	0
4	0	0	0	0	1	0	0	0
5	0	0	0	1	0	0	0	0
6	0	0	1	0	0	0	0	0
7	0	1	0	0	0	0	0	0
8	1	0	0	0	0	0	0	0
Streaths								
1	1	1	1	1	1	1	1	0
2	1	1	1	1	1	1	0	1
3	1	1	1	1	1	0	1	1
4	1	1	1	1	0	1	1	1
5	1	1	1	0	1	1	1	1
6	1	1	0	1	1	1	1	1
7	1	0	1	1	1	1	1	1
8	0	1	1	1	1	1	1	1

Note. The 4-d conditions contained either Dimensions 1–4 or Dimensions 5–8. The 8-d conditions included all eight dimensions. The 1s and 0s refer to values on a dimension (D).

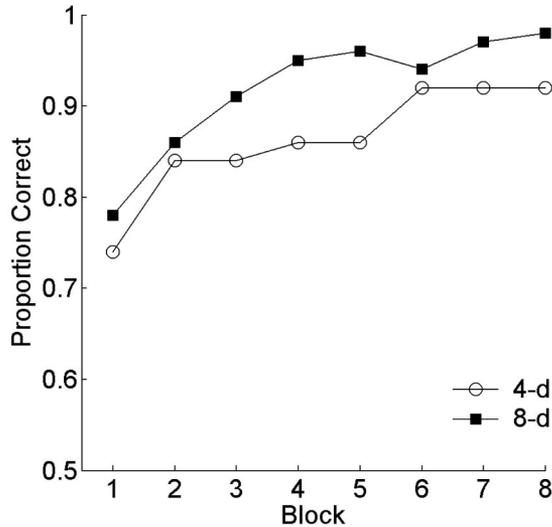


Figure 1. Proportion correct as a function of learning block and category structure in Experiment 3.

Discussion

The 8-d advantage present in Experiments 1 and 2 for the single-feature items was replicated. Because we controlled the number of learning blocks, as in Experiment 2, we can attribute the advantage to the dimensionality variable rather than to differences in exposure to the stimuli. In contrast to our previous learning results, which were consistent with associationist accounts *with* limited attention, the results here are consistent only with associationist learning accounts *without* limited attention. Furthermore, the 8-d advantage for single-feature items was even more notable in the current experiment compared with the previous two, because it occurred with an increased learning rate for the 8-d compared with the 4-d condition. Thus, considering together the data from the learning and transfer phases, our results are inconsistent with standard accounts of category learning. They are consistent, however, with the unlabored speed by which children and adults learn real-world categories—categories that have a large number of dimensions instead of a few.

General Discussion

The impetus for our studies was the apparent disparity between the ease of learning natural categories with many dimensions on the one hand and the difficulty caused by additional dimensions in experimental categories on the other. To this end, we examined category learning rates and knowledge of the individual dimensions for categories that differed in their dimensionalities but were the same in their structural coherence ratios. Contrary to our expectations from current accounts of categorization, but consistent with how people learn rich natural categories, we found that increasing the number of stimulus dimensions did not harm learning of individual features. In fact, subjects in the 8-d condition were able to learn many of the additional features provided, even when learning more features was not necessary for correct performance. Neither did more dimensions lead consistently to worse category learning: In Experiment 3, the 8-d condition even exhib-

ited significantly higher classification accuracy than the 4-d condition over the eight learning blocks. Our learning results were consistent with those of Minda and Smith (2001, Experiment 4), who found performance to be very similar for the learning of four- and eight-dimensional stimuli, but who did not test individual-feature knowledge.

Learning of features is an important measure, because it is the basis for much category use. When, for example, one hears that a friend has a new puppy, one infers that the animal likely barks, chews slippers, needs to be house-trained, and so on. The use of categories in induction, communication, and comprehension relies largely on knowing what features different categories have. This is why we have focused on this measure in our study, rather than, say, learning speed, which has been of interest in other studies.

This raises another issue, namely, the strategy by which people learned the categories. No doubt, different people in our experiments used different strategies, such as rule-testing and prototype or exemplar learning (as proposed by, e.g., Malt, 1989; Smith & Minda, 1998; and see next section), and it is possible that the different structures we used encouraged different learning strategies. Unfortunately, there is no well-accepted empirical measure to determine which strategies individual subjects use (at least, no measure based on classification data; cf. Rehder & Hoffman, 2005a). Although it would be revealing to have information about such strategies, we believe that the issue of feature learning cuts across these different strategies. That is, as we just pointed out, whether people are learning exemplars or testing rules, they eventually have to learn that dogs bark, chew bones, play fetch, and so on, or else they will not be able to identify, reason, or talk about dogs—which are the functions of categories. We believe, then, that this issue is one that must be addressed by every theory of category learning, although most past studies have focused on classification of entire objects (see Markman & Ross, 2003, for a detailed discussion). We have made our discussion as general as possible, so that it will be relevant to whatever strategy of category learning people use. Nevertheless, our results may in fact have implications for different theories, which we discuss next.

Implications for Category-Learning Models

Hypothesis-testing theories make the clearest predictions about the effect of dimensionality on learning. Recall that having more dimensions increases the size of the rule space (Markman & Maddox, 2003; Markman & Ross, 2003). Simply stated, with more dimensions, there are more rules to test, and in our categories in particular, if learners tested all one- or two-dimensional rules before arriving at a correct three-dimensional rule, they would have taken much longer to learn the higher-dimensional categories. That did not in fact happen. Indeed, our learning measures indicated that either the 8-d and 4-d category structures were of about equal difficulty (Experiments 1 and 2) or the 8-d category structure was actually easier to learn (Experiment 3). Furthermore, equating the number of dimensions necessary to learn the categories did not result in equal amounts learned about the categories (in Experiment 3; see discussion below), contrary to predictions of the hypothesis-testing approach. Thus, it is clear that our subjects were not using pure hypothesis testing for the learning of the two category structures.

In discussing other theories, we need to consider how they would predict our single-feature categorization data, which have not generally been used to test different models. Exemplar models could explain single-feature categorization by analogy to classification of entire exemplars, with the presented feature activating exemplars possessing that feature. Exemplar models based on the Medin and Schaffer (1978) context model have the simplifying assumption that an exemplar is perfectly encoded once seen. Performance in feature classification would then depend on parameters for similarity and decision computations (i.e., the c parameter, which represents exemplar sensitivity, and γ , which represents the level of deterministic responding). If we assume that such parameters are constant across 4-d and 8-d conditions (and at a reasonable range of values), exemplar models would clearly be consistent with learning more dimensions in the 8-d case. However, the assumption that each exemplar is immediately and perfectly encoded into memory is problematic. In many past experiments, this assumption did no harm, because the exemplars did not differ systematically in their memorability. But in the present paradigm, it seems unrealistic to assume that people can learn 4-d and 8-d exemplars equally easily. Thus, further development of the exemplar approach to this paradigm seems necessary. Exemplar models also often incorporate weight parameters that limit the amount of attention to each dimension. We discuss such attentional issues in a later section.

Multiple-system theories of categorization suggest that people can use a number of learning strategies simultaneously. Thus, besides hypothesis testing, people may memorize exemplars (Erickson & Kruschke, 1998; Nosofsky et al., 1994) or form a decision boundary using multiple stimulus dimensions (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby & Waldron, 1999). For example, both ATRIUM (Erickson & Kruschke, 1998) and COVIS (Ashby et al., 1998) incorporate a competitive gating mechanism that selects between single-dimension and multiple-dimension strategies based on their relative performance. These models could learn our categories by forming a rule on one dimension and then learning exception exemplars (ATRIUM) or integrate the other dimensions (COVIS) to account for the rule's shortcoming (as explained earlier, our categories would require a three-dimensional rule, which these models do not currently use). This strategy would be equally successful for our 4-d and 8-d structures, although one might question whether learning the exemplars made up of eight features ought to be as easy as learning those made of four.

The implications of our results for hypothesis testing and exemplar memorization, discussed above, apply to multiple-system models to the extent that one or the other system gets deployed. To the extent that a model classifies with rules, it should predict slower learning of the 8-d than the 4-d category structure and equal learning of features in the two structures in Experiment 3 (where three-dimensional rules led to perfect classification for both structures). On the other hand, if the exemplar-memorization component is used, then the model needs to address the perfect encoding assumption discussed earlier. Of course, when different strategies can be mixed within a single learner, predicting a precise pattern of results can be very difficult, as the two components can interact to produce behavior that neither on its own would produce. Thus, we cannot say for certain what any mixed model would do when

applied to our task. Our data provide an interesting new test for such models.

As discussed in the introduction, whereas simple associative network models predict a category-learning advantage for the 8-d condition, which we did not consistently find, associative networks with limited attention can explain no change in learning rate with added dimensions, as attention limitations can trade off the associative benefit. Thus, our learning results are generally consistent with such models. However, these models have a problem with the feature-learning data, as attentional limitations should cause less learning about each dimension, as the limited attentional pool is spread over more dimensions. Furthermore, error-driven learning algorithms seem to predict that these models should stop learning after acquiring the minimum number of dimensions to allow perfect performance.

Error-driven learning and attentional limitations have considerable evidence supporting them in the category-learning literature, and we should be careful in concluding that they cannot account for our results. In fact, we suspect that the problem is not that they are wrong but that other factors may mitigate their effects in the learning situations studied here. We discuss each principle in detail in the next sections.

Error-Driven Learning

One reason that learning the properties of more complex categories might be harder is that not all dimensions are necessary for performance. After one has reached a high level of accuracy in classifying exemplars based on a subset of the dimensions, there should be little change in the association weights of features (or exemplars), because there is little error to drive the change in the weights. (Systems incorporate this feature, presumably, because when performance is accurate, weight changes cannot improve accuracy but could reduce it.) Thus, this principle predicts that adding more stimulus dimensions will not result in further dimensions being learned if they do not promote greater accuracy, just as the second cue in blocking or overshadowing is not learned in classical conditioning.

Experiments 1 and 2 could be consistent with this principle, because some subjects probably had to learn more stimulus dimensions in the 8-d condition in order to reach criterion. If any two of the dimensions they learned were redundant, then they would have had to acquire three more dimensions in order to accurately classify the category members. Thus, the finding that subjects learned more dimensions in the 8-d than in the 4-d condition could be accommodated by this assumption. (However, it is not clear that every error-driven learning algorithm could predict a greater number of dimensions learned in the 8-d condition given that there were not actually more errors in that condition to drive more learning.)

However, things are not as positive for error-driven learning in Experiment 3, because here the number of dimensions necessary to reach a given level of accuracy was identical across the two conditions. Thus, according to this principle, once subjects correctly learned three dimensions, learning should have stopped in both conditions, because there would have been no further error. But instead of learning equivalent numbers of dimensions, subjects in the 8-d condition learned more. And in fact, they made consis-

tently fewer errors during learning. This, then, is a puzzle for such learning mechanisms.

Limited Attention

The other principle we mentioned as predicting a problem in learning larger categories is the assumption that attention is limited. In particular, most models since Nosofsky's (1984) assume that weights on dimensions must sum to 1, and therefore as more dimensions are added, the weight given to any dimension on average must decrease. Our results seem inconsistent with this prediction, because we found learning of more dimensions in the larger categories and responses to individual features that were generally just as fast as those of people who had learned fewer dimensions.

There are a number of possible responses to this argument, related to the fact that attention weights are not directly reflected in performance. One response is to suggest that in fact attention was reduced across dimensions in the 8-d condition, but because the amount of attention necessary to learn dimensions was low, even this reduced attention sufficed. Imagine, for example, that a minimum attention weight of .10 is necessary to learn a dimension and that attention weights must sum to 1.0. In the 4-d condition, each dimension has .25 on average, and so its dimensions can be learned. (Not all dimensions will necessarily get the minimum and be learned, however.) In the 8-d condition, each dimension has a .125 weight on average, so each one may also be learned. Thus, even though attention to each dimension was reduced by half in the 8-d condition, more stimulus dimensions were learned, because only a little attention was necessary for each dimension, and even the 8-d categories had many dimensions that exceeded the minimum.

This explanation is not consistent with the RT results, however, as it clearly predicts that the weight associated with each dimension is lower in the 8-d condition (half as great, on average). Thus, the individual dimensions are not as strongly associated with their category in the 8-d condition. Yet we did not find a reliable RT increase for the 8-d group in the tests of individual features. (In Experiment 1, there was a difference only if a potentially outlying subject was removed, but in Experiment 2, there was a marginal speed advantage for the 8-d condition.) Furthermore, our strict criterion, requiring correct responses on all four questions about a stimulus dimension, also showed more dimensions learned in the 8-d case. If attention weights in the 8-d condition were half those (on average) in the 4-d condition, it is difficult to understand how those subjects could reach this high criterion more often. Lower weights should have led to more variable responding and hence to more errors.

A different response of the attentional spreading explanation is to suggest that the observed results are due to a statistical artifact, in which the number of dimensions learned was truncated in the 4-d condition but not the 8-d condition. For example, suppose that our subjects could learn on average six dimensions. Because no one in the 4-d condition could learn more than the presented four but those in the 8-d condition could learn up to eight, the performance of the 4-d group, but not the 8-d group, was artificially truncated.

There is some validity to this point in principle, but it cannot account for the difference between conditions. In particular, the

truncation explanation predicts that the number of people who learned four dimensions in the 4-d condition should be the same as the number who learned four or more in the 8-d condition. In fact, across all three experiments the number of subjects learning all four dimensions in the 4-d condition was only 14 out of 48 (29%)—considerably fewer than the number in the 8-d condition who learned four or more dimensions, 34 of 48 (71%) ($p < .01$, Fisher's exact test, two-tailed). This comparison leads to the surprising conclusion that something in the situation of having more stimulus dimensions makes one learn *more*, rather than reducing learning. We discuss this surprising implication below.

Finally, note that attention weights and error-driven learning explanations have been provided for learning situations with many fewer dimensions. For example, the typical case of blocking in classical conditioning involves two different stimuli: Learning one conditioned response inhibits the learning of a second, redundant one. In probability learning experiments, it is common to test two stimulus dimensions (Kruschke & Johansen, 1999). Thus, it is not consistent to use these principles to explain why it is that learning an initial stimulus dimension interferes with learning a second and at the same time to suggest that learning eight dimensions in our category-learning task does not incur significant attentional limitations or yield learning competition among dimensions. We find the past explanations of interference in learning to be compelling and are not questioning them. What we are questioning is whether these explanations apply in the present situation, using family resemblance categories in a category-learning task.

Differences Between Category Learning and Other Learning Tasks

In an attempt to come to grips with the surprising results we have reported, we believe that it would be helpful to point out some differences between the situation we have investigated—family resemblance category learning—and the situations most often used as evidence for limited attention and error-based learning, the probability learning task and classical conditioning. Of course, such a discussion at this stage must be somewhat speculative.

Consider first classical conditioning. In Kamin's (1969) blocking paradigm, he first conditioned rats in the *blocking condition* to expect a shock after, say, a flash of light. In the second stage, all rats (blocking and control) received a shock after a compound stimulus consisting of the light plus a tone. Then, to assess association strength, Kamin measured how much the hungry rat slowed its eating in response to one of the stimuli. Whereas the control rats slowed eating in response to the tone, the rats in the blocking condition did not. In other words, the rats did not learn that the tone from the compound stimulus predicted a shock if they had already learned that the light predicted the shock. Learning the light-shock association blocked learning of the redundant tone-shock association.

Some of the differences between our experimental paradigm and the conditioning paradigm are so obvious that they hardly bear mentioning. Yet it may be those differences that are the important ones. Clearly, the rat is not learning a concept. The light and tone are physically unconnected events that do not seem to arise from a single object or class of objects. The rat is not interested in learning about a concept—it only wishes to recognize when the

shock will come. So, the rat does not need to learn additional information. We might contrast this with a situation in which the rat encounters a novel animal, with a variety of visual properties, sounds, actions, smells, and so forth. Now the features are unified in a single object rather than being truly independent cues, and the rat might now learn more than one property of this animal. Our human subjects had a situation more like this, in contrast to the unrelated light and tone in conditioning experiments. Thus, the goals inherent in a learning situation may be a critical aspect to consider.

Similarly, in the probability learning situation studied by Edgell et al. (1992), among others, subjects learn to predict one of two outcomes, which are simply two different responses (e.g., left or right key). The two cues are chosen to be perceptually independent. For example, Edgell (1978) used Xs and Ys that were colored red or green. Because of the probabilistic nature of the rules used, subjects could not reach perfect accuracy. Indeed, Kruschke and Johansen (1999) told their subjects to try to reach 70%–80% accuracy.

In some sense, then, these subjects are also not learning categories. That is, they are not being presented with two different kinds of objects that are in different equivalent classes. In fact, the exact same object has different correct responses on different trials of the experiment, unlike almost every natural category. (We know of no instance in which identical animals are in two different species, for example.) Furthermore, there may be no intrinsic structure to the stimuli, in that all combinations of stimulus features can occur equally often (as in Kruschke & Johansen's [1999] Figure 3). If all possible combinations of features are equally likely in the stimulus set, it is difficult to specify categories in the sense of structure in the environment that people can learn and exploit (Rosch, 1978).

Intuitively, then, we would like to suggest that subjects in the classical conditioning and probability learning experiments are not learning categories and therefore may not be acting like people who are trying to learn categories. What do true category learners do that is different?

One difference may be that subjects in our experiments took as their task the goal to learn as much about the categories as possible, whereas subjects in the other experiments simply wished to make correct responses. The fact that four or eight dimensions were presented in our stimuli carries an implication that each of these dimensions is relevant to being a Mobble or a Streath: Mobbles don't only have distinctive eyes and wings, they also have their own body marking and feet. One might perform accurately having only learned three dimensions, but in the 8-d condition, one would not then be learning what Mobbles are really like. Therefore, even after learning has reached a fairly high level (as in Experiment 3, when performance was greater in the 8-d condition throughout learning), subjects with high-dimensional categories continue to learn more dimensions. Perhaps presenting more stimulus dimensions actually engages additional attention, such that people attempt to learn more. This assumption seems necessary to explain why people learned more dimensions in the 8-d condition than in the 4-d condition: Something must have spurred them to attempt to learn more, even when it was not strictly necessary.

This explanation relies on the fact that we studied family resemblance categories in which the different stimulus dimensions were all predictive of category membership. We did not use

nonlinearly separable or weakly structured categories with non-predictive dimensions. Thus, the presented dimensions all provided information about the category if the subjects wished to learn more about it. This is in contrast to many popular experimental category structures, in which some stimulus dimensions are uninformative (see Murphy, 2005). Clearly, if learning about extra dimensions is difficult or not informative, learners might not put forth sufficient effort to do so. An important goal for future work on this topic is to extend the research to different category structures.

Finally, we might compare our results with those of mainstream category-learning experiments, which have strongly suggested that people give more or less attention to different dimensions, based on their relevance to classification (e.g., Medin & Schaffer, 1978; Nosofsky, 1986; Rehder & Hoffman, 2005b; Smith & Minda, 2000). Doesn't this strongly established result provide convincing evidence that attention is limited and therefore that dimensions compete? We suspect that the difference is that in these experiments, giving attention to some dimensions is actually harmful to category learning, because they are misleading or very weak predictors of category membership. That was not the case here.

That said, we should note that the term *attention* is used somewhat indirectly in all of these discussions. In most studies, attention weights are inferred from classification. Thus, attention weights are a measure of how much a stimulus dimension is *used in classification*. However, that is not the same as how much a person has *learned about* that stimulus dimension. If a learner has realized that values on a dimension do not predict category membership, then this learner must have attended to this dimension at one point and has in fact learned about it. However, in modeling, this person would yield a low weight to the dimension, because he or she would not use it in classification. Thus, one must be careful not to equate low weights on a dimension with having learned little about that dimension. (Consistent with this notion, Kruschke & Johansen, 1999, used the term *cue utilization* rather than *attention*.) For this reason, we believe that our results are not in principle incompatible with past modeling efforts that have assumed a fixed amount of "attention" across dimensions. That attention reflects the decision weight placed on each dimension, whereas our post-test measured what subjects learned about how each dimension relates to category membership. Each construct may be important to understanding how categories are learned and classification accomplished. But one must distinguish the form of attention involved in learning from the possibly separate issue of decision weights.

Models of Associative Learning

We have suggested that the limitations of learning and attention proposed by current theories of category learning do not apply in a straightforward way to family resemblance category learning of the sort studied in our experiments. In particular, we have proposed that when people take their task as being category learning, they recruit the attention necessary to learn as many of the stimulus dimensions as possible. Although we believe that the principles of error-driven learning and attention limitations certainly do not predict these results, the results seem more consistent with one of these accounts than the other.

Error-driven learning has a great difficulty with our results. The reason is that accounts such as the Rescorla–Wagner learning rule have limits on what can be learned about a particular outcome (category, in our case; unconditioned stimulus in the conditioning case) (Pearce & Bouton, 2001). Put simply, once you are able to predict the outcome, you will learn no more about what predicts it. This is presumably a basic limit on associative learning processes and so does not seem to be something that could be overcome by strategies or goals, which we have proposed as explaining our results.

In contrast, attentional explanations of the same phenomena seem more easily extendable to our results. According to Mackintosh's (1975) explanation, successful stimuli gain attention strength during learning. Once the outcome is successfully predicted, little attention is given to new cues, and so they are not learned. In a different approach, Pearce and Hall (1980) proposed that when an outcome is correctly predicted, little attention is given to *any* cue on subsequent trials (not just new ones). Thus, although more could be learned about the outcome, in the blocking situations it is not, because of this decrease in attention. Note that this attentional limitation is not the same as the one used in many categorization models, in that it does not specify a fixed pool of attention that is spread over dimensions. Instead, the total amount of attention to cues decreases as a result of (successful) learning, regardless of the number of dimensions.

The Pearce and Hall (1980) view in particular seems susceptible to alteration to account for the present results. As is well known, attention is to some degree controllable through executive processes that direct the allocation of resources (Baddeley, 1986). Thus, if people normally do not attend to cues after learning is achieved, it should be possible to overcome this tendency by the reallocation of attention. In our experiments, learners are no doubt thinking something like "I wonder which kind of bug has two eyes and which has four eyes?" even after they have started to classify the bugs fairly well. Because Pearce and Hall did not claim that the long-term status of a cue is changed by the learning process, they need not predict that later presented cues will not be learned if something draws attention to them. We are arguing that a reallocation of attention to unlearned dimensions allows subjects in the 8-d condition to continue to learn after the category can be correctly predicted. In contrast, the Rescorla–Wagner model claims that there is nothing to be learned about the response, and so its entire explanatory mechanism would have to be changed in order to account for further learning.

Implications for Future Research

As we remarked, our explanation of our findings is speculative. The focus of the present article has been to establish that, in fact, categories with more dimensions do not suffer interference such that number of learned dimensions stays the same or is reduced. This result, which is rather surprising in light of the history of studies showing interference among learning cues, is a significant one. However, it must be understood that we have looked at only one kind of category structure (simple family resemblance designs) and one stimulus domain (visually defined bugs).³ One important consideration may be what happens when some of the stimulus dimensions are not predictive, because that had been the cause for lowering attention to some dimensions in past research.

More generally, our explanation predicts that if categories are not very "category-like," these effects may not obtain. If the stimulus dimensions are clearly independent cues, not part of a single object, and if the outcome is an arbitrary response, then people may not engage in the attention shifting necessary to learn multiple dimensions. Instead, they may be perfectly happy with learning the minimal dimensions necessary for good performance.

Finally, we return to the apparent contradiction we noted between natural category learning, in which learning many dimensions is normal, and classical conditioning and probabilistic cue learning, in which learning multiple dimensions sometimes does not occur. It would be surprising if the basic learning processes themselves were different in these two situations. Instead, we have suggested that task demands and strategies account for this difference. In the simple cue-learning situation, prediction of the outcome is the goal, and the traditional learning models likely explain why there is this limitation on how many cues are learned. In category-learning situations with naturalistic categories, learners make an effort to learn as much about the category as possible—not just to predict the category correctly. Indeed, watching children at a zoo or reading a picture book reveals their interest in acquiring as much information as possible about (at least some) categories. They could stop looking at the zebra or squirrel once they've learned about its stripes or tail, yet they are apparently motivated to discover and learn much more. Although such intrinsic motivation may not be as strong or universal in the college student population learning experimental categories, we suspect that this or a similar variable is responsible for why such learners seem to acquire more information about a category than they need in order to complete the task.

³ However, we have found the 8-d feature-learning advantage using textual stimuli and family resemblance categories in a recent study.

References

- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.
- Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, *30*, 666–677.
- Ashby, F. G., & Waldron, E. M. (1999). On the nature of implicit categorization. *Psychonomic Bulletin & Review*, *6*, 363–378.
- Baddeley, A. (1986). *Working memory*. New York: Oxford University Press.
- Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. Oxford, England: Wiley.
- Carey, S. (1978). The child as word learner. In M. Halle, J. Bresnan, & G. A. Miller (Eds.), *Linguistic theory and psychological reality* (pp. 264–293). Cambridge, MA: MIT Press.
- Edgell, S. E. (1978). Configural information processing in two-cue non-metric multiple-cue probability learning. *Organizational Behavior and Human Performance*, *22*, 404–416.
- Edgell, S. E., Bright, R. D., Ng, P. C., Noonan, T. K., & Ford, L. A. (1992). The effect of representation of the processing of probabilistic information. In B. Burns (Ed.), *Percepts, concepts and categories: The repre-*

- sentation and processing of information. *Advances in psychology* (Vol. 93, pp. 569–601). Amsterdam: Elsevier.
- Edgell, S. E., Castellan, N. J., Roe, R. M., Barnes, J. M., Ng, P. C., Bright, R. D., et al. (1996). Irrelevant information in probabilistic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1463–1481.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, 127, 107–140.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- Heit, E. (1994). Model of the effects of prior knowledge on category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1264–1282.
- Judd, C. M., & McClelland, G. H. (1989). Interactions and polynomial regression: Interactions between predictor variables. In J. Kagan (Ed.), *Data analysis: A model-comparison approach* (p. 247–264). San Diego, CA: Harcourt Brace Jovanovich.
- Kamin, L. J. (1969). Predictability surprise, attention, and conditioning. In R. Church & B. A. Cambell (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Kruschke, J. K. (1993). Three principles for models of category learning. In G. V. Nakamura, R. Taraban, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 29): *Categorization by humans and machines* (pp. 57–90). San Diego, CA: Academic Press.
- Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1083–1119.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276–298.
- Maddox, W. T. (1995). Base-rate effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 288–301.
- Malt, B. C. (1989). An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 539–555.
- Malt, B. C., & Smith, E. E. (1982). The role of familiarity in determining typicality. *Memory & Cognition*, 10, 69–75.
- Markman, A. B., & Maddox, W. T. (2003). Classification of exemplars with single- and multiple-feature manifestations: The effects of relevant dimension variation and category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 107–111.
- Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, 129, 592–613.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 355–368.
- Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, 19, 242–279.
- Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 775–799.
- Murphy, G. L. (2001). Fast-mapping children vs. slow-mapping adults: Assumptions about words and concepts in two literatures. *Behavioral and Brain Sciences*, 24, 1112–1113.
- Murphy, G. L. (2003). Ecological validity and the study of concepts. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 43, pp. 1–41). San Diego, CA: Academic Press.
- Murphy, G. L. (2005). The study of concepts inside and outside the lab: Medin vs. Medin. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin* (pp. 179–195). Washington, DC: American Psychological Association.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology*, 52, 111–139.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian conditioning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, 87, 532–552.
- Rehder, B., & Hoffman, A. B. (2005a). Eyetracking and selective attention in category learning. *Cognitive Psychology*, 51, 1–41.
- Rehder, B., & Hoffman, A. B. (2005b). Thirty-something categorization results explained: Selective attention, eyetracking, and models of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 811–829.
- Rehder, B., & Murphy, G. L. (2003). A knowledge-resonance (KRES) model of category learning. *Psychological Bulletin & Review*, 10, 759–784.
- Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4, 328–350.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Rosch, E. H., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Rosch, E., Mervis, C., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382–439.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, 75(13, Whole No. 517).
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1411–1436.
- Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 3–27.
- Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 3–11.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113, 169–193.
- Widrow, G., & Hoff, M. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention, Convention Record*, 4, 96–104.

Received February 12, 2005

Revision received August 1, 2005

Accepted October 11, 2005 ■