

What's Behind Different Kinds of Kinds: Effects of Statistical Density on Learning and Representation of Categories

Heidi Kloos
University of Cincinnati

Vladimir M. Sloutsky
The Ohio State University

This research examined how differences in category structure affect category learning and category representation across points of development. The authors specifically focused on category density—or the proportion of category-relevant variance to the total variance. Results of Experiments 1–3 showed a clear dissociation between dense and sparse categories: Whereas dense categories were readily learned without supervision, learning of sparse categories required supervision. There were also developmental differences in how statistical density affected category representation. Although children represented both dense and sparse categories on the basis of the overall similarity (Experiment 4A), adults represented dense categories on the basis of similarity and represented sparse categories on the basis of the inclusion rule (Experiment 4B). The results support the notion that statistical structure interacts with the learning regime in their effects on category learning. In addition, these results elucidate important developmental differences in how categories are represented, which presents interesting challenges for theories of categorization.

Keywords: categorization, category learning, selective attention, cognitive development

The ability to form categories (i.e., treating discriminable entities as members of an equivalence class) is a critically important component of human cognition. There is much research on categorization and category learning spanning early infancy to adulthood (see Murphy, 2002, for a review). Some results of this research point to an interesting paradox: Whereas some categories are learned in an effortless and unsupervised manner even by young infants (Eimas & Quinn, 1994; Quinn, Eimas, & Rosenkrantz, 1993; Younger & Cohen, 1986), other categories are difficult even for adult learners who are given feedback after each trial (e.g., Bruner, Goodnow, & Austin, 1956). In this article, we discuss a solution to this paradox, focusing on the role of category structure on category learning and category representation across points of development.

The Paradox of Category Learning

Consider the two sides of the above mentioned paradox. On the one hand, young infants can exhibit effortless unsupervised category learning of even ill-defined categories such as *cat* and *dog* (Eimas & Quinn, 1994; Mareschal & Quinn, 2001; Quinn et al., 1993). For example, in Quinn et al. (1993), 3- to 4-month-olds were first familiarized with members of the category (e.g., pictures of cats) and then shown new members of the studied category (e.g.,

new cats) paired with members of a nonstudied category (e.g., dogs). If infants learn the category during familiarization, they should discriminate new members of studied categories from members of nonstudied categories. Results indicated that with just six familiarization trials, participants exhibited evidence of category learning.

On the other hand, adults can have difficulty learning categories that are well defined by Boolean algebra rules (Bruner et al., 1956; Shepard, Hovland, & Jenkins, 1961). For example, in some of the Bruner et al. (1956) experiments, participants had to learn a category that included all and only items with green circles. On each trial, participants were asked to determine whether a particular item belonged to the target category, and their responses were followed by feedback. Despite the fact that these were well-defined, strictly deterministic categories and despite the fact that each trial was accompanied by feedback, learning of some categories elicited substantial difficulties.

Thus, some well-defined categories are difficult to learn even for adult learners, whereas some ill-defined categories are effortlessly learned even by 3- to 4-month-olds. The ease of category learning shown by infants cannot stem from the use of familiar categories given that it was also demonstrated with novel artificial stimuli (e.g., Bomba & Siqueland, 1983; Younger, 1993; Younger & Cohen, 1986). Nor could the difficulty shown by adults stem from the use of category-inclusion rules that are more complicated than those used in infant studies. Animal categories are ill defined and probabilistic and thus have more complex inclusion rules than the deterministic well-defined categories used in the Bruner et al. (1956) study. Finally, it is unlikely that the learning paradox stems from a difference in stimulus complexity given that the infant studies demonstrating category learning used both perceptually rich stimuli (Quinn et al., 1993) and perceptually impoverished ones (Bomba & Siqueland, 1983).

Heidi Kloos, Department of Psychology, University of Cincinnati; Vladimir M. Sloutsky, Center for Cognitive Science, The Ohio State University.

This research has been supported by National Science Foundation Grant REC 0208103 to Vladimir M. Sloutsky. We thank Anna Fisher, Chris Robinson, Brad Love, and Aaron Yarlas for their helpful feedback.

Correspondence concerning this article should be addressed to Vladimir M. Sloutsky, Center for Cognitive Science, 208C Ohio Stadium East, 1961 Tuttle Park Place, The Ohio State University, Columbus, OH 43210. E-mail: Sloutsky.1@osu.edu

We suggest that the explanation of the paradox has to do with the difference in category structure between the categories used in the infant studies and the categories used in adult studies. This research focuses on an important aspect of category structure—the *statistical density* of the category, or the ratio of category-relevant variance to the total variance. More specifically, we suggest that category structure interacts with the learning regime: Dense categories can be learned without supervision, whereas learning of sparse categories requires supervision.

Before we further define this concept and its relation to learning regime, it is important to note that statistical density is not the only measure of category structure (for a review, see Medin, Lynch, & Solomon, 2000). Structural differences considered before pertain to syntactic differences (i.e., nouns vs. verbs; e.g., Gentner, 1981), ontological differences (i.e., natural kinds vs. nominal kinds; e.g., Kripke, 1972), taxonomic differences (i.e., basic level vs. superordinate level; e.g., Rosch & Mervis, 1975), differences in content (i.e., entity categories vs. relational categories; e.g., Gentner & Kurtz, 2005), differences in concreteness (i.e., concrete vs. abstract categories; e.g., Barsalou, 1999), difference in linear separability (i.e., linearly separable vs. nonlinearly separable; e.g., Medin & Schwanenflugel, 1981; Waldron & Ashby, 2001), difference in category coherence and confusability (e.g., Homa, Rhoades, & Chambliss, 1979; Rouder & Ratcliff, 2004; J. D. Smith & Minda, 2000), differences in redundancy (e.g., Garner, 1962), and category utility (Corter & Gluck, 1992). However, as we argue in the General Discussion, below, statistical density may map onto some of these distinctions—while offering important advantages.

Statistical Density

Any set of items can have a number of possible dimensions (e.g., color, shape, size), some of which might vary and some of which might not. Categories that are statistically dense have multiple intercorrelated features relevant for category membership, with only a few features being irrelevant. Good examples of statistically dense categories are basic-level animal categories such as *dog*. Dogs have a particular range of shapes, sizes, and colors; they have four legs and a tail; and they bark. These features are jointly predictive, thus yielding a dense (albeit probabilistic) category.

Categories that are statistically sparse have very few common features, with the rest of the features varying independently and thus constituting a set of irrelevant or surface features. Good examples of sparse categories are scientific concepts such as *accelerated motion*. Consider two events: (a) a planet revolving around a sun and (b) a cat chasing a mouse. Only a single relation—the change in the planet’s and the cat’s vector of motion—makes both events variants of accelerated motion. All other features and feature relations are irrelevant for membership in this category, and they can vary greatly.

Conceptually, statistical density is a ratio of variance relevant for category membership to the total variance across members and nonmembers of the category. Therefore, density is a measure of statistical redundancy (Shannon & Weaver, 1948), which is an inverse function of relative entropy. In general, density is a measure of nonrandomness or regularity, whereas entropy is a measure of randomness. The advantage of expressing category structure though entropy is that entropy has a great deal of computational

plausibility given that living organisms are claimed to be able to automatically detect entropy in a set (e.g., Young & Wasserman, 2001).

Density can be expressed as

$$D = 1 - \frac{H_{within}}{H_{between}}, \quad (1)$$

where H_{within} is the entropy observed within the target category and $H_{between}$ is the entropy observed between target and contrasting categories. In what follows, we explain statistical density in greater detail. Two aspects of stimuli are important for calculating statistical density, variation in stimulus dimensions and variation in relations among dimensions.

First, stimulus dimension may vary either within a category (e.g., members of a target category are either black or white) or between categories (e.g., all members of a target category are black, whereas all members of a contrasting category are white). Within-category variance decreases density, whereas between-category variance increases density. We make a simplifying assumption that the varying dimensions are binary (e.g., the size of an entity is either big or small).

Second, dimensions of variation may be related (e.g., all items are black circles), or they may vary independently of each other (e.g., items can be black circles, black squares, white circles, or white squares). Covarying dimensions result in smaller entropy than dimensions that vary independently. On the basis of previous evidence (cf. Whitman & Garner, 1962), we assume that only dyadic relations (i.e., relations between two dimensions) are detected spontaneously, whereas relations of higher arity (e.g., a relation among color, shape, and size) are not. Therefore, only dyadic relations are included in the calculation of entropy.

The total entropy is the sum of the entropy due to varying dimensions (H^{dim}) and the entropy due to varying relations among the dimensions (H^{rel}). More specifically,

$$H_{within} = H_{within}^{dim} + H_{within}^{rel}, \quad (2a)$$

and

$$H_{between} = H_{between}^{dim} + H_{between}^{rel}. \quad (2b)$$

The concept of entropy was formalized by information theory (Shannon & Weaver, 1948), and we use these formalisms here. First, consider the entropy due to dimensions. This within-category and between-category entropy is presented in Equations 3a and 3b respectively.¹

$$H_{within}^{dim} = - \sum_{i=1}^M w_i \left[\sum_{j=0,1} \text{within}(p_j \log_2 p_j) \right], \quad (3a)$$

and

$$H_{between}^{dim} = - \sum_{i=1}^M w_i \left[\sum_{j=0,1} \text{between}(p_j \log_2 p_j) \right], \quad (3b)$$

¹ Note that whenever a value of a dimension (or a relation between dimensions) has a probability of zero, these values are not included into the calculation because they do not contribute to the entropy.

where M is the total number of varying dimensions, w_i is the attentional weight of a particular dimension (the sum of attentional weights equals a constant), and p_j is the probability of value j on dimension i (e.g., the probability of a color being white). The probabilities could be calculated within a category or between categories.

The attentional-weight parameter is of critical importance—without this parameter, it would be impossible to account for learning of sparse categories. In particular, when a category is dense, even relatively small attentional weights of individual dimensions add up across many dimensions. This makes it possible to learn the category without supervision. Conversely, when a category is sparse, only few dimensions are relevant. Supervision is therefore necessary to direct attention to these relevant dimensions.

Note that a maximal entropy due to a single dimension is observed when two values of a dimension are equally probable (e.g., half of the items are white and half of the items are black). Assuming that the attentional weight of the dimension $w_i = 1.0$, the maximal entropy due to this dimension is 1. Any deviation from equal probability of each value reduces the maximal entropy.

Next, consider the entropy that is due to a relation between dimensions. To express this entropy, we need to consider the co-occurrences of dimensional values. If dimensions are binary, with each value coded as 0 or 1 (e.g., white = 0, black = 1, circle = 0, and square = 1), then the following four co-occurrence outcomes are possible: 00 (i.e., white circle), 01 (i.e., white square), 10 (i.e., black circle), and 11 (i.e., black square). The within-category and between-category entropy that is due to relations is presented in Equations 4a and 4b, respectively.

$$H_{within}^{rel} = - \sum_{k=1}^O w_k \left[\sum_{\substack{m=0,1 \\ n=0,1}} p_{mn} \log_2 p_{mn} \right], \quad (4a)$$

and

$$H_{between}^{rel} = - \sum_{k=1}^O w_k \left[\sum_{\substack{m=0,1 \\ n=0,1}} p_{mn} \log_2 p_{mn} \right], \quad (4b)$$

where O is the total number of possible dyadic relations among the varying dimensions, w_k is the attentional weight of a particular relation (again, the sum of attentional weights equals a constant), and p_{mn} is the probability of a co-occurrence of values m and n on dimension k . Again, these probabilities could be calculated either within a category or between categories. As shown above, when values are binary, mn can take values of 01, 01, 10, and 11. Given that the total number of varying dimensions is M , the number of dyadic relations O can be calculated using Equation 5:

$$O = \frac{M!}{(M-2)! * 2!}.$$

Note again that the entropy is maximal when each outcome is equally probable, that is, when dimensions vary independently of each other and each outcome occurs with the probability of .25. If the attentional weight of a relation $w_k = 1.0$, the entropy due to a single relation is 2. However in reality, the attentional weight of a relation is likely to be less than 1.0 given that a relation is more

difficult to detect than a single dimension. On the basis of empirical data, we estimate that the weight of a relation is no more than half of the weight of a dimension (see Appendix A for supporting empirical evidence). Therefore, if we make a simplifying assumption that the attentional weight of a dimension is 1.0, then the attentional weight of a relation is 0.5.

For example, suppose that there are two dimensions of variation (e.g., color and shape). All entities in the target category are white circles, whereas all entities in the contrasting category are black squares. Therefore, the within-category probability of white circles, white squares, black circles, and black squares is 1.00, 0, 0, and 0, respectively, whereas the between-category probability of these feature pairs is .50, 0, 0, and .50, respectively. The within-category entropy due to the two dimensions $H^{dim} = 0 + 0$, and the within-category entropy due to the color–shape relation $H^{rel} = 0$. Therefore, the total within-category entropy $H_{within} = 0$, yielding a category density of $D = 1.00$.

Now suppose that the relation between the dimensions is weaker. Entities in the target category are white circles and black squares, whereas entities in the contrasting category are black circles and white squares. The within-category probability of white circles, white squares, black circles, and black squares is .50, 0, 0, and .50, respectively, whereas the between-category probability of these feature pairs is .25, .25, .25, and .25, respectively. In this case, nonweighted within-category entropy due to the two dimensions is $1 + 1 = 2$ (weighted $H^{dim} = 2.0$), and nonweighted within-category entropy due to the relation is 1 (weighted $H^{rel} = 0.5$). The total weighted within-category entropy $H_{within} = 2.5$, whereas the total weighted between-category entropy $H_{between} = 3.0$, yielding a category density of $D = 0.17$.

In what follows, we apply calculations of statistical density to the well-known categories used by Shepard et al. (1961). Stimuli in that study could vary on the three binary dimensions of size, color, and shape (see Figure 1). We assume the attentional weight of each dimension $w^{dim} = 1.0$, the attentional weight of each relation $w^{rel} = 0.5$. Given that the same stimuli were used for all six types of categories, the between-category entropy is the same. Specifically, weighted between-category entropy due to dimensions $H^{dim} = 3.0$, and weighted between-category entropy due to relations $H^{rel} = 3.0$. This yields $H_{between} = 6.0$.

In the Type I category, color is a single fully predictive dimension (i.e., all target items are black and all contrasting items are white), with shape and size being equally distributed between target and contrasting items. Therefore, the within-category entropy due to dimensions is 2. The nonweighted within-category entropy due to the relation color–size is 1 ($p_{black\ large} = .50$, $p_{black\ small} = .50$, $p_{white\ large} = 0$, and $p_{white\ small} = 0$). The same is true for the nonweighted within-category entropy due to the relation color–shape ($p_{black\ square} = .50$, $p_{black\ triangle} = .50$, $p_{white\ square} = 0$, and $p_{white\ triangle} = 0$). The nonweighted within-category entropy due to the relation shape–size is 2 ($p_{large\ square} = .25$, $p_{large\ triangle} = .25$, $p_{small\ square} = .25$, and $p_{small\ triangle} = .25$). Therefore, the weighted within-category entropy due to relations $H^{rel} = 2.0$. Taken together, the total within-category entropy $H_{within} = 4.0$, and the resulting density of this category $D = 1 - (4/6) = 0.33$.

In the Type II category, the only predictive feature is the relation between color and shape. Target items are white squares and black triangles, and contrasting items are black squares and white triangles. Given that individual dimen-

Category Type:	Type I	Type II	Type III	Type IV	Type V	Type VI
Target Items	■ ▲ ■ ▲	□ □ ▲ ▲	■ ▲ ▲ △	■ ▲ ▲ △	■ ▲ ▲ □	▲ △ □ ■
Contrasting Items	□ □ △ △	■ ■ △ △	□ △ □ ■	□ △ □ ■	□ △ △ ■	△ ▲ ■ □
Density	0.33	0.08	0.19	0.23	0.12	0

Figure 1. Category types used in the Shepard, Hovland, and Jenkins (1961) study: target items, contrasting items, and statistical densities.

sions are not predictive, the within-category dimensional entropy $H^{dim} = 3.0$. The nonweighted within-category entropy due to the predictive relation color–shape is 1 ($p_{white\ square} = .50$, $p_{black\ triangle} = .50$, $p_{black\ square} = 0$, and $p_{white\ triangle} = 0$), and the nonweighted within-category entropy due to each of the non-predictive relations of shape–size and color–size is 2 ($p_{large\ square} = p_{large\ triangle} = p_{small\ square} = p_{small\ triangle} = p_{large\ white} = p_{large\ black} = p_{small\ white} = p_{small\ black} = .25$). Therefore, the weighted relational entropy $H^{rel} = 2.5$, the total weighted within-category entropy $H_{within} = 5.5$, and the resulting density of this category $D = 1 - (5.5/6) = 0.08$. Densities of other category types used by Shepard et al. (1961) are presented in Figure 1.

Category Density and Category Learning

There is much evidence that sets of items with intercorrelated dimensions are learned better than sets of items with uncorrelated dimensions (Billman & Knutson, 1996; Garner, 1962; Whitman & Garner, 1962). For example, participants exhibited better recall of sets of geometric shapes that had simple within-set contingencies (i.e., greater category density) than of those that did not (Whitman & Garner, 1962). However, it is also known that people can ably learn sparse categories defined by a single dimension, such as categories defined only by color or only by shape (e.g., Kruschke, 1992; Nosofsky, 1986; Shepard et al., 1961; Trabasso & Bower, 1968; see also L. B. Smith, 1989, for a developmental proposal). Interestingly, most of the studies demonstrating an advantage of dense categories typically used unsupervised learning paradigms, whereas studies demonstrating the ability to learn sparse categories mainly used supervised learning paradigms.

We argue that the reason for this disparity has to do with selective attention. Dense categories put small demands on selective attention because they have a high ratio of relevant to irrelevant information, and as a result, learning of statistically dense categories does not require supervision. In fact, supervision may hinder learning of dense categories. For example, an explicit description of the many relevant features and feature correlations in a dense category might put a high memory demand on the learner and therefore make category learning more difficult. Su-

pervision may also invite participants to look for rules, something that might impede learning of a dense category (cf. Reber, 1976).

Conversely, statistically sparse categories put high demands on selective attention because they have a low ratio of relevant to irrelevant information. The learner not only needs to know what to focus on but also needs to ignore the large proportion of irrelevant information. As a result, learning of statistically sparse categories is likely to require some form of top-down information that specifies which dimensions are to be attended to and which are to be ignored. Such top-down information could involve various kinds of external supervision, including explicit instruction, corrective feedback, guided comparisons, or negative evidence. The most direct way of communicating what is relevant might be to provide the learner with the category-inclusion rule.

If these considerations are correct, then it is reasonable to expect that category structure would interact with the learning regime. In particular, dense categories could be ably learned without supervision, whereas sparse categories would require supervision. Furthermore, if learning of statistically sparse categories puts high demands on selective attention, then it is reasonable to expect developmental differences in learning of sparse categories. This is because early in development, the ability to selectively attend to relevant information is less pronounced than later in development (Kirkham, Cruess, & Diamond, 2003; Napolitano & Sloutsky, 2004; Zelazo, Frye, & Rapus, 1996; see also Dempster & Corkill, 1999, for a review).

Category Density and Category Representation

Density of a category may affect not only category learning but also category representation. Traditionally, researchers considered two types of representations (see Murphy, 2002, for a review). One type is a rule-based representation (i.e., a category is represented by its inclusion rule), and the other type is a similarity-based representation (i.e., a category is represented by either a prototype or a set of exemplars). More recently, several hybrid models that include both rule- and similarity-based representations have been proposed (e.g., RULEX: Nosofsky, Palmeri, & McKinley, 1994; ATRIUM: Erickson & Kruschke, 1998; COVIS: Ashby, Alfonso-

Reese, Turken, & Waldron, 1998; see also E. E. Smith & Sloman, 1994). Does statistical density affect the way a category is represented? If yes, how does it affect category representation?

We consider several possibilities. First, it is possible that both dense and sparse categories are represented in a rule-based manner (cf. Bruner et al., 1956). For example, even the dense category *cat* may be represented by the rule “everything that meows.” Alternatively, it is possible that both dense and sparse categories are represented in a similarity-based manner (cf. Allen & Brooks, 1991). For example, even the sparse category *accelerated motion* may be represented by examples of what such movement looks like. Yet another alternative is that dense and sparse categories are represented differently. For dense categories, the learner might form similarity-based representations given that appearance features are likely to be relevant for category membership. For sparse categories, on the other hand, the learner might form rule-based representations given that appearance features are likely to be irrelevant for category membership. Finally, it is possible that the type of category representation changes in the course of development. Similarity-based representations may appear early in development (French, Mareschal, Mermillod, & Quinn, 2004), whereas rule-based representations may develop later. This last possibility is consistent with the recently proposed similarity-based theory of early categorization (e.g., Sloutsky, 2003; Sloutsky & Fisher, 2004), as well as with some earlier-proposed theoretical views (e.g., see Keil, 1989).

Overview of Current Experiments

The goal of the current experiments was to test the hypotheses formulated above and to examine how dense and sparse categories are learned and represented across points of development. To achieve this goal, we conducted a series of experiments with adults and preschool children in which we manipulated the statistical density of the category. Experiments 1–3 focused on category learning, whereas Experiment 4 focused on category representation.

In Experiments 1–2, we examined how dense and sparse categories are learned in adults and children. Participants were presented with dense or sparse categories under an unsupervised learning condition (learners were presented only with members of the target category²) or under a supervised learning condition (learners were given the explicit description of the category-inclusion rule). The principal difference between the two experiments was that categories were linearly separable in Experiment 1 (i.e., the categories could be separated on the basis of summed dimensional values), whereas categories were not separable in Experiment 2 (i.e., the target and contrast categories were defined by relations among dimensions and not by individual dimensions).

Experiment 3 examined whether statistical density is a better predictor of category learning than alternative predictors, such as the total number of dimensions or the absolute number of relevant dimensions. Finally, in Experiment 4, we examined the representation of dense and sparse categories in adults and children. This time, dense and sparse categories were acquired under the same learning condition: Participants were presented with an explicit description of the target category and with individual members of the target category. After training, participants were given a surprise recognition task in which they had to determine whether a

test item had been presented during training. If participants form a similarity-based representation of the learned category, they should false-alarm on new items that are similar in appearance to training items. On the other hand, if participants form a rule-based representation of the learned category, they should false-alarm on new items that have the same rule as the training items.

Experiment 1

The goal of Experiment 1 was to examine the effects of category density on category learning in adults (Experiment 1A) and in children (Experiment 1B). The experiment used artificial creature-like stimuli that had several dimensions of variation (e.g., the shading of the body, the size of the wings, the number of antennas, etc.). In the statistically dense category, all of these dimensions were predictive of category membership (i.e., items of the target category had a dark body, long wings, etc.). In the statistically sparse category, only one of the dimensions (e.g., the shading of the body) was predictive, whereas all other dimensions varied randomly. Category information was presented either in an unsupervised learning condition (participants were merely presented with members of the category) or in a supervised learning condition (participants were explicitly told the category-inclusion rule). On the basis of the considerations presented above, we expected that the statistically sparse category would require supervision, whereas the statistically dense category could be ably learned without supervision.

Experiment 1A

Method

Participants. For this and subsequent experiments with adults, participants were Introductory Psychology students at The Ohio State University who participated in the experiment for a partial course credit. Sixty participants took part in this experiment (41 women and 19 men). Additionally, 2 participants were tested and omitted from the sample because their accuracy on catch trials did not meet the criterion (see *Procedure*, below).

Materials and design. The stimuli were colorful drawings of artificial creatures (see Figure 2). Each instance could vary on six dimensions: size of tail, size of wings, number of buttons, number of fingers, shading of body, and shading of antennas. Each dimension had two levels (e.g., a short wing and a long wing). However, to ensure some variability among category members, each level had two values (e.g., the short wing could be 1.8 or 2.3 cm, and the long wing could be 4.6 or 5.6 cm).

Stimuli were created such that they belonged either to a statistically dense category (i.e., each dimension was predictive of the category membership) or to a statistically sparse category (i.e., only one dimension was predictive). Table 1 shows the structure of the target and contrasting items in abstract notation. Target items of the dense category had short tails, short wings, few buttons, few fingers, light bodies, and light antennas, whereas contrasting items had long tails, long wings, many buttons, many fingers, dark bodies, and dark antennas. For the sparse category, target items

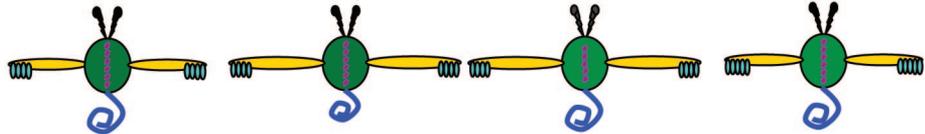
² The term *unsupervised* refers to a manner of learning that lacks any top-down information about the category's inclusion rule.

Dense Category

Target Items

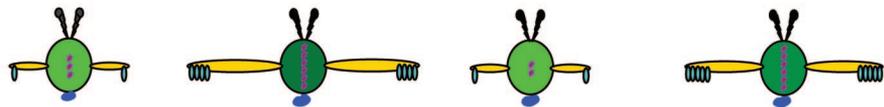


Contrasting Items



Sparse Category

Target Items



Contrasting Items

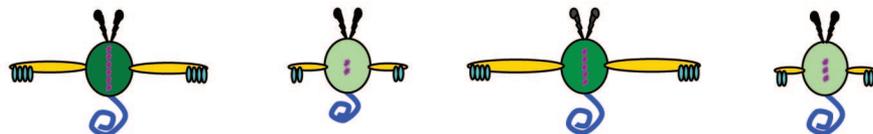


Figure 2. Examples of stimuli used in Experiment 1.

had a short tail, and contrasting items had a long tail. All other dimensions varied randomly across items. The statistical density of the dense category was 1.00, and the density of the sparse category was 0.17 (for detailed calculations, see Appendix B).

The experiment had a 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised) between-subjects design. Participants were randomly assigned to one of the four resulting conditions.

Procedure. In this and all other experiments with adults reported here, participants were tested in a quiet room on

campus. The experiments were administered on a computer and were controlled by SuperLab Pro 2.0 software (Cedrus Corporation, San Pedro, CA). Participants were instructed that they had to learn to distinguish fictitious creatures called Ziblets (i.e., target items) from creatures that were not Ziblets (i.e., contrasting items).

The procedure included a training phase and a testing phase, with only the training phase differing across learning conditions. During training in the unsupervised learning condition, participants were presented with 16 target instances, one by

Table 1
Structure of Stimuli Used in Experiment 1

Dimension	Dense category		Sparse category	
	Target item	Contrast item	Target item	Contrast item
Size of tail	0	1	0	1
Size of wings	0	1
Number of fingers	0	1
Number of buttons	0	1
Shading of body	0	1
Shading of antennas	0	1

Note. The numbers 0 and 1 refer to the features of the respective dimension (e.g., 0 = short, 1 = long). Features varied randomly in the cells marked with an ellipsis (. . .). Each feature had two levels within a category to allow for some variability among items.

one, in a self-paced manner. During training in the supervised learning condition, participants were not presented with any instances. Instead, they were presented with a statement describing the necessary and sufficient features of target items. For the dense category, participants were given the following statement: “Ziblets have a light body, a short tail, two or three buttons, short yellow wings, and one or two fingers on each yellow wing.” For the sparse category, the statement was, for example, “Ziblets have a short tail.” Each feature mentioned in the statement about the sparse category was accompanied by a picture showing the isolated feature. Note that participants did not obtain any information about the contrasting category in either learning condition.

The testing phase was administered immediately after the training phase. Participants were presented with 32 test items, half of which were target items and half of which were contrasting items. The task was to determine whether or not the shown item was a Ziblet. Eight catch trials followed as a test of a participant’s overall alertness. These catch trials consisted of non-Ziblets with new features (i.e., they had a diamond-shaped body, triangle-shaped wings, and no tail). It was expected that participants accurately reject these items regardless of category structure or learning condition. To be included in the study, participants had to reject at least six out of the eight catch items.

Results and Discussion

To evaluate learning of the target category, accuracy scores were calculated for each participant across the 32 test items. Accuracy scores represented the difference between the proportions of hits (i.e., correct identification of a Ziblet) and false alarms (i.e., an incorrect identification of a non-Ziblet as a Ziblet). In principle, these scores can vary from 1 (i.e., perfect discrimination between Ziblets and non-Ziblets) to -1 (i.e., a “reverse discrimination” between Ziblets and non-Ziblets). However, in practice, the reverse discrimination and resulting negative scores are unlikely. It was therefore expected that scores would vary from 0 (i.e., no discrimination) to 1 (i.e., perfect discrimination).

Mean accuracy scores by category type and learning condition are presented in Figure 3A. These accuracy scores were subjected to a 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised) between-subjects analysis of variance (ANOVA). The analysis revealed a significant interaction, $F(1, 56) = 19.60, p < .001$. Although mean accuracy was above zero in all conditions (single-sample t s $> 8.10, p < .001$), the dense category was learned better in the unsupervised than the supervised condition, independent-sample $t(30) = 4.30, p < .001$, whereas the sparse category was learned better in the supervised condition than the unsupervised condition, independent-sample $t(26) = 3.30, p < .01$. These results support the hypothesis that category structure interacts with learning regime.

Having found the interaction between category density and category learning in adults, we deemed it necessary to examine this interaction early in development. If learning of sparse categories puts greater emphasis on selective attention than learning of dense categories, then it is reasonable to expect that young children would exhibit greater difficulty than adults in acquiring statistically sparse categories without supervision. At the same time,

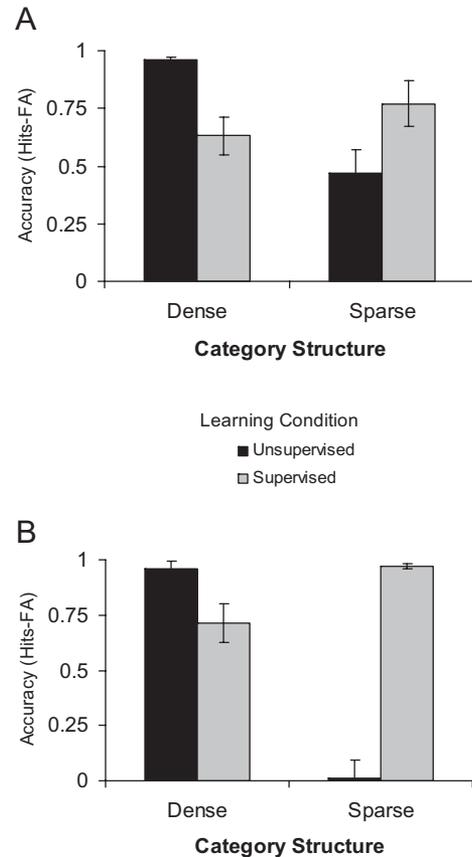


Figure 3. Mean accuracy scores by category type and learning condition in Experiment 1A (Panel A) and 1B (Panel B). Error bars represent standard errors of the mean. FA = false alarms.

because acquisition of dense categories does not put emphasis on selective attention, young children should ably acquire these categories without supervision.

Experiment 1B

Method

Participants. For this and all subsequent experiments with children, participants were 4- and 5-year-olds recruited from pre-schools located in middle-class suburbs of Columbus, Ohio. Participants in this experiment were 49 children ($M_{age} = 58.6$ months, $SD = 3.7$ months; 25 girls and 24 boys) assigned randomly to one of the four conditions used in Experiment 1A.

Materials, design, and procedure. Materials, design, and procedure were similar to those in Experiment 1A, with the following exceptions. First, children were given a cover story involving a character who would like to get a pet from a magical store. Second, each rule (for the dense and sparse categories) was repeated three times, and each feature description was again accompanied by a picture showing the isolated feature. Finally, to shorten the experiment, catch trials were not used in the testing phase.

Results and Discussion

Mean accuracy scores by category type and learning condition are presented in Figure 3B. A 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised) between-subjects ANOVA rendered the predicted interaction significant, $F(1, 45) = 62.00, p < .001$. For the dense category, average accuracy scores were significantly higher in the unsupervised than the supervised condition, $t(21) = 2.28, p < .05$, whereas for the sparse category, the scores were significantly higher in the supervised than in the unsupervised condition, $t(24) = 10.91, p < .01$. Therefore, the dissociation between dense and sparse categories was observed again: Children ably learned the dense category without supervision, whereas learning of the sparse category required supervision.

Furthermore, in contrast with adults, who exhibited evidence of learning of the sparse category in the unsupervised condition (see Experiment 1A), children exhibited no evidence of learning the sparse category in the unsupervised condition (single-sample $t < 1.0$). This finding suggests that young children, who have difficulty in deliberately controlling their attention, could not spontaneously discover the relevant dimension, focus on this dimension, and ignore irrelevant dimensions in the course of category learning.

Experiment 1C

Experiments 1A and 1B present evidence that participants successfully learned dense categories without supervision, whereas learning of sparse categories required supervision, at least for young children. However, these experiments leave an important question unanswered. Did participants rely on the holistic pattern of correlated features when learning a dense category without supervision, or did they focus on a single feature? If the latter possibility is the case, it can have one of two variants. First, it is possible that most participants focused on the same single feature, with learning of this feature being supported by the presence of other correlated features. Alternatively, it is possible that different participants focused on different features—in the dense condition, many features are predictive of category membership, and focusing on any of them would result in successful category learning (cf. Trabasso & Bower, 1968). To distinguish among these possibilities, testing items were modified in such a way that the correlated structure of features was lost.

As was done in Experiments 1A and 1B, adults and children were asked to learn the dense category through the unsupervised learning condition, and their learning was then assessed by asking them to distinguish unmodified target items (Ziblets) from unmodified contrasting items (Flurps). An additional testing phase was added that contained a crucial manipulation: Ziblets had a planted feature of a Flurp, and Flurps had a planted feature of a Ziblet. The important measure was the degree to which participants categorized the modified items on the basis of the planted feature. One of four different features was planted in modified stimuli. To account for the possibility that different participants might focus on a different feature when categorizing an item, we calculated for a participant four separate categorization scores, one categorization score for each

of the four individual features. We then compared a participant's highest feature-based categorization score of modified items with his or her categorization score of unmodified items.

If items are categorized on the basis of the single feature, then a participant's highest categorization score of the modified items should be comparable with his or her categorization score of unmodified stimuli. However, if participants attempt to categorize on the basis of the holistic pattern of correlated features, then categorization scores of modified items should be substantially lower than those of unmodified items.

Method

Participants. Participants were 18 children ($M_{age} = 58.5$ months, $SD = 3.3$ months; 10 girls and 8 boys) and 22 adults (11 women and 12 men). None of them participated in the previous experiments.

Materials, design, and procedure. To determine whether learners pay attention to the pattern of correlated features of the dense category (rather than extract an isolated feature during the learning task), we created a second set of testing stimuli for which a salient feature of one category (e.g., the tail of a Ziblet) was planted into the stimulus of the other category (resulting in a modified Flurp). Each of these features was planted one at a time, thus resulting in four types of modified stimuli: those that had a planted wing, those that had a planted tail, those that had a planted body, and those that had planted antennae.³

The procedure consisted of a training phase and two testing phases. The training phase was identical to the training phase used in Experiments 1A and 1B for unsupervised learning of the dense category: Participants were shown target stimuli, one by one, in a self-paced manner, and asked to learn about Ziblets. The ensuing testing phases required participants to categorize stimuli presented for a short time (200 ms for adults and 750 ms for children). This speed pressure was introduced to elicit the most basic pattern of categorization and to prevent participants from deploying multiple categorization strategies.⁴ During the first testing phase, learners had to categorize unmodified items (half of which were Ziblets and half of which were Flurps). This was done to establish how well participants had learned the dense category. During the second testing phase, learners had to categorize the modified stimuli (half of which were Ziblets that had one Flurp feature and half of which were Flurps that had one Ziblet feature).

If participants base their categorization of dense categories on individual features, the following prediction can be made: For each participant, there should be at least one type of modified stimuli for which feature-based categorization of these items is comparable to categorization of unmodified items. This should not be the case, however, if participants categorize dense categories on the basis of

³ Pilot testing revealed that wings, tail, antennae, and body of the creature represented the most salient features of the stimuli. In this pilot experiment, adult participants ($n = 16$) were presented with target and contrasting items of the dense category and asked to write down features differentiating Ziblets and Flurps. The four features were mentioned at least once by at least one participant.

⁴ These presentation times were determined in a separate calibration experiment as sufficient for distinguishing between unmodified target and contrasting item by children and adults.

the entire pattern. Here, categorization of unmodified items should exceed feature-based categorization of modified items.

Results and Discussion

Five categorization scores were calculated for each participant, one to reflect a participant's categorization of unmodified items and four to reflect a participant's categorization of modified items—one for each type of planted feature (antenna, body, tail, and wing). Each categorization score could vary from 0 to 1.0, with .5 reflecting chance performance. The categorization score for unmodified items measured the proportion of correctly categorizing unmodified Ziblets as Ziblets and unmodified Flurps as Flurps. The categorization score for modified items reflected the proportion of categorizing on the basis of the planted feature (e.g., categorizing a Ziblet with Flurp's tail as a Flurp).

The primary analysis focused on participants' categorization scores of unmodified versus modified items. Only one categorization score of modified items was used for each participant, namely, the one that was highest for the participant. This is because different participants could have relied on different individual features, in which case averaging across subsets of

modified items would have resulted in a loss of this information. For example, if a participant categorizes on the basis of a creature's tail, the only relevant feature-based categorization score would be the categorization score of stimuli that have a planted tail. Another participant might focus on a creature's body, in which case this participant's highest feature-based categorization score would be for the subset of stimuli that have a planted body. We then compared this highest feature-based categorization score of modified items with a participant's categorization of unmodified items.

The results are presented in Figure 4A. Data in the figure clearly indicate that categorization of unmodified items exceeded feature-based categorization of modified items for both children and adults. A two-way mixed ANOVA ($\text{Age} \times \text{Stimulus Type}$) indeed revealed a significant main effect of stimulus type, $F(1, 39) = 34.08, p < .0001$. Neither the main effect of age nor the interaction approached significance (both $ps > .25$). These findings indicate that categorization of unmodified items was based on the entire pattern rather than on a single feature, thus undermining the possibility that participants categorized dense categories on the basis of a single feature.

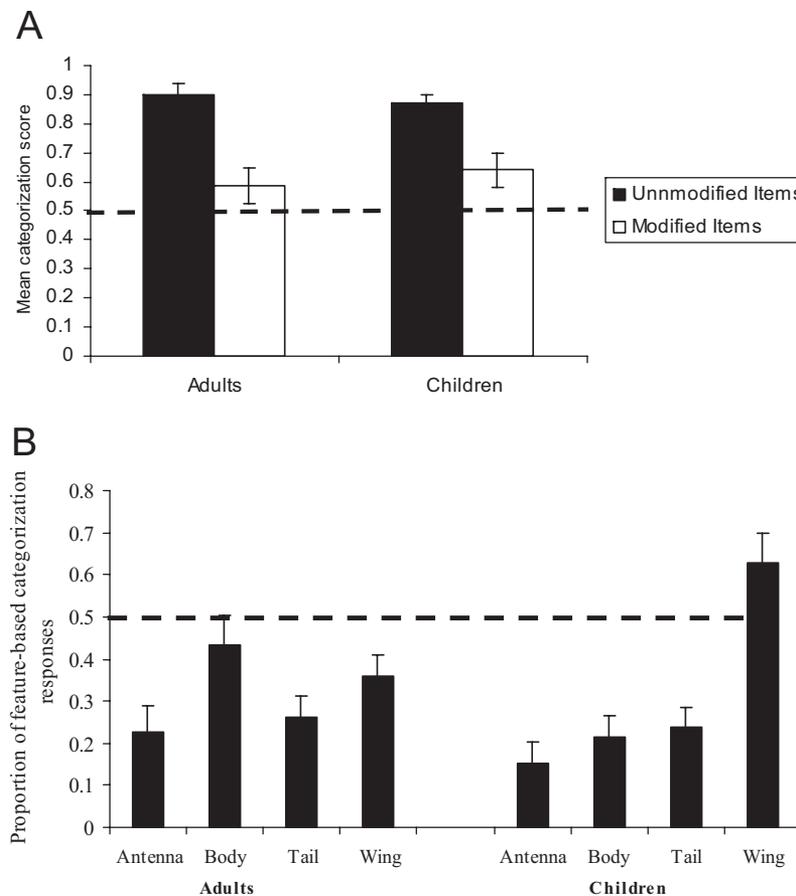


Figure 4. Mean categorization scores in Experiment 1C for adults and children. Error bars represent standard errors of the mean. A: Categorization score of unmodified stimuli and best feature-based categorization score of unmodified stimuli. B: Feature-based categorization scores by the type of modified stimuli (antenna, body, tail, and wing).

Could it be that participants were merely tired because modified items were presented in the second block, with fatigue generating chance performance? If this is the case, then reliance on all planted features should be around chance. At the same time, if participants tended to focus on the overall pattern even when items were modified, then reliance on the remaining planted features should be below chance. The results of feature-based categorization responses for the four planted features are presented in Figure 4B. They indicate that for three out of four planted features, both children and adults exhibited below-chance feature-based categorization (all one-sample $t_s > 3.80$, $p_s < .002$).

Taken together, these results strongly indicate that unsupervised learning of dense categories is driven by the entire pattern of correlated features rather than by attending to a single feature. This is true even though the dimensions are separable, not integral. These results provide further evidence that dense and sparse categories are learned differently: Although learners focused on the isolated predictive feature when learning the sparse category, they focused on the pattern of correlated features when learning the dense category.

Experiment 2

One could argue that efficient learning of the statistically dense category shown in Experiment 1 stemmed from high within-category similarity rather than from high statistical density. High density in Experiment 1 was confounded with high similarity among category members, but high density does not have to co-occur with high similarity. A more stringent test of statistical density would be to use statistically dense categories that do not have high within-category similarity (i.e., categories that are not linearly separable).

To achieve this goal, we modified stimuli from Experiment 1 in such a way that category membership was predicted by a relation between dimensions rather than by one or more individual dimensions. As a result, the same features occurred in both the target and the contrasting categories, which substantially attenuated the within-category similarity. Several relations were predictive of category membership for the statistically dense category, but they varied independently for the statistically sparse category (with only one relation being predictive of category membership). Finding a dissociation between dense and sparse categories using nonlinearly separable categories would expand the generalizability

of findings of Experiment 1. This hypothesis was tested with adults (Experiment 2A) and children (Experiment 2B).

Experiment 2A

Method

Participants. Participants were 60 adults (28 women and 32 men), none of whom participated in the previous experiments, who were randomly assigned to one of the four conditions used in Experiment 1. Additionally, nine adults (between two and three in each condition) were tested and excluded from the sample because their performance on the catch items did not reach the criterion (see *Procedure*).

Materials and design. The stimuli were similar to those used in Experiment 1A with one important difference: Category membership was determined by relations among dimensions rather than by individual dimensions. For the statistically dense category (density = 0.39; see Appendix B for details of density calculations), all dimensions covaried within the category. For example, an instance with a short tail, short wings, and long fingers had a dark body, dark antennae, light buttons, few fingers, and few buttons. Examples of stimuli are shown in Figure 5, and their abstract structure is presented in Table 2. Note that category members differed from each other in terms of individual features. For example, one member could have a short tail, and another member could have a long tail. Furthermore, target items did not differ from contrasting items in terms of individual features. For example, both target and contrasting items could have a short tail (see Figure 5).

To ensure that within-category similarity was indeed comparable to between-category similarity, we asked a separate group of adults ($n = 19$) to rate the similarity between target items (i.e., within-category similarity) and the similarity between target and contrasting items (i.e., between-category similarity). A 9-point rating scale was used, with 1 representing *not similar at all* and 9 representing *very similar*. The results point to virtually equivalent within- and between-category similarity, with the mean within-category similarity being 5.27 and the mean between-category similarity being 5.00, paired sample $t(18) < 1.00$.

For the statistically sparse category, the varying dimensions were the size of wings, the shading of antennae and body, and the number of tails, buttons, and fingers. Only one relation was cate-

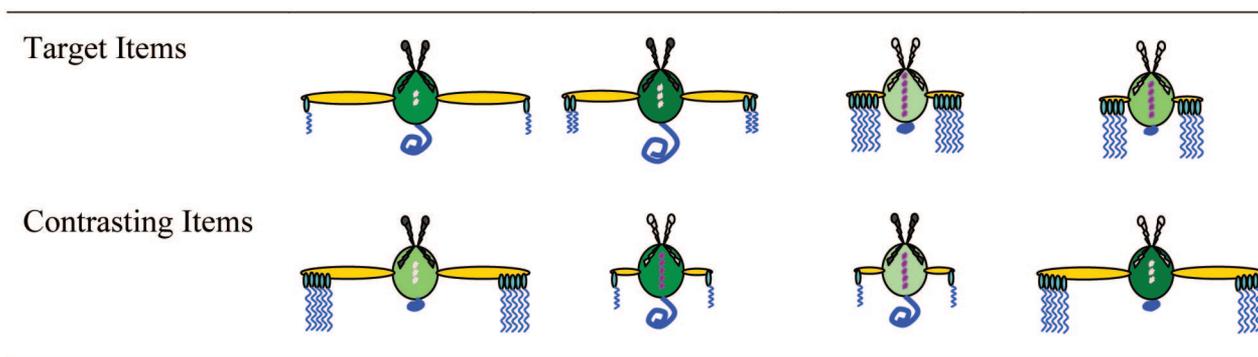


Figure 5. Examples of stimuli for the dense-category conditions used in Experiment 2.

Table 2
Structure of Stimuli in the Dense Category Used in Experiment 2

Dimension	Target category		Contrasting category	
	Item 1	Item 2	Item 1	Item 2
Size of tail	0	1	0	1
Size of wings	0	1	1	0
Size of fingers	1	0	0	1
Shading of body	0	1	0	1
Shading of antennas	0	1	1	0
Shading of buttons	1	0	0	1
Number of fingers	0	1	0	1
Number of buttons	0	1	1	0

Note. The numbers 0 and 1 refer to the features of the respective dimension (e.g., 0 = short tail, 1 = long tail). Each feature had two levels to allow for some variability among items.

gory relevant, with the other dimensions varying randomly. More specifically, target items had fewer body buttons than tails plus fingers, whereas contrasting items had more body buttons than tails plus fingers. For example, a target item could have five buttons, three tails, and four fingers, resulting in fewer buttons than tails plus fingers. The numbers of buttons, tails, and fingers were chosen in such a way that neither the number of a single feature nor the correlation between two of the features was predictive. This ensured that no other information (e.g., difference in quantity) was redundant with the inclusion rule. Given that the inclusion rule was based on a triadic relation that was not expected to be detected spontaneously, the total density of the sparse category was 0. Therefore, this category should have been difficult, if not impossible, to learn without supervision.

As in Experiment 1, this experiment had a 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised) between-subjects design. Participants were randomly assigned to one of the four resulting conditions.

Procedure. The number of training trials in the unsupervised learning condition was increased to 32 to ascertain that participants would learn the more complex nonlinearly separable categories. The explicitly stated inclusion rule for the statistically dense category was “A Ziblet with a dark body has dark antennas, long wings, a long tail, one or two short fingers, and two or three light buttons; and a Ziblet with a light body has light antennas, short wings, a short tail, four or five long fingers, and five or six dark buttons.” For the statistically sparse category, the explicitly stated inclusion rule was “For a Ziblet, the number of buttons is smaller than the number of tails and fingers together.” To make the mathematical relation of the sparse category clear, the rule was accompanied by an example in which a particular number of buttons, fingers, and tails were depicted separately.

Results and Discussion

Mean accuracy scores (the proportion of hits minus the proportion of false alarms) by category type and learning condition are presented in Figure 6A. Category learning was again a function of statistical density and learning condition. A 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised)

between-subjects ANOVA confirmed the significant interaction, $F(1, 56) = 46.14, p < .001$. As in Experiment 1, participants learned the statistically dense category better without than with supervision, whereas they learned the statistically sparse category better with supervision than without (independent-sample t s $> 3.30, ps < .001$). Furthermore (and unlike Experiment 1), the sparse category used in Experiment 2 was not learned at all without supervision, with mean accuracy not being different from zero, $t(14) = 0.41, p > .68$.

These results replicate and further extend findings of Experiment 1 by using statistically dense yet nonlinearly separable cat-

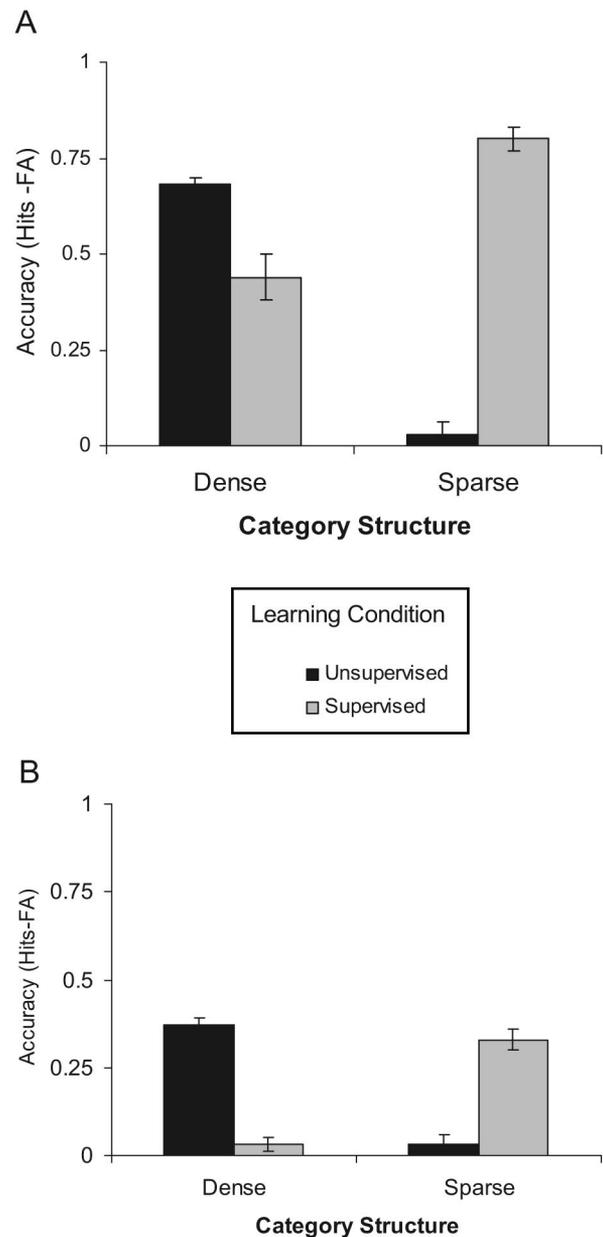


Figure 6. Mean accuracy scores by category type and learning condition in Experiments 2A and 2B. Error bars represent standard errors of the mean. FA = false alarms.

egories. The results suggest that the dissociation between category structure and learning regime is driven by the difference in statistical density rather than by a difference in within- versus between-category similarity.

Although the idea of statistical density can account for the entire pattern of results, it is possible that other factors might account for results of individual cells. First, it is possible that learning of dense categories in the supervised learning condition was affected by the length and difficulty of the category's inclusion rule. In particular, participants might have thought that they had to pay attention to all the features mentioned in the rule. To address this possibility, a new group of adults ($n = 15$; 7 women and 8 men) was tested in their ability to learn the dense category through explicit instruction. The only difference in the procedure was that the verbal description of the dense category was shorter, containing only two, rather than six, features. For example, adults were told, "A Ziblet with a dark body has dark antennas, and a Ziblet with a light body has a short tail." Despite this simplification, the mean accuracy in the ensuing test ($M = 0.63$, $SE = 0.07$) did not exceed the accuracy score of participants who learned the dense category through mere exposure ($M = 0.68$, $SE = 0.05$). This finding supports the conclusion that explicit instruction about the inclusion rule does not facilitate learning of a dense category even when the inclusion rule is shorter.

Second, it is possible that adults could not learn the sparse category through the unsupervised learning regime because the sparse category was based on an unusually difficult category-inclusion rule. To address this issue, we constructed a sparse category that was comparable to the dense category in that the correlation between two features mattered. Specifically, the correlation between a creature's number of fingers and number of buttons was predictive of category membership, whereas all other feature correlation varied randomly. In abstract notation, the structure of target items was 11xxxxxx and 00xxxxxx (Ziblets with many fingers had many buttons, and Ziblets with few fingers had few buttons), whereas the structure of the contrasting items was 10xxxx and 01xxxx (non-Ziblets with many fingers had many buttons, and non-Ziblets with few fingers had few buttons). The density of this category was 0.03 (see Appendix B). A new group of adults ($n = 24$; 13 women and 11 men) learned this category either in the unsupervised learning regime (they were presented with members of the target category) or in the supervised learning regime (they were given the explicit inclusion rule of the target category). The results show the predicted effect of learning regime, $t(22) = 4.42$, $p < .01$, with the mean accuracy score being higher under the supervised learning regime ($M = 0.58$, $SE = 0.06$) than under the unsupervised learning regime ($M = 0.11$, $SE = 0.07$). Thus, the effect of learning regime on the learning of a sparse category was confirmed even with a more straightforward and simpler sparse category.

Experiment 2B

The goal of Experiment 2B was to examine learning of nonlinearly separable categories early in development. If category density contributes to category learning, then it is reasonable to expect that even young children should be able to learn a nonlinearly separable dense category without supervision. Furthermore, it appears highly unlikely that young children are capable of learning a sparse

category without supervision given that adults did not exhibit such learning. Finally, given the difficulty of ignoring irrelevant information early in development, it was expected that young children would exhibit weaker learning of this exceedingly sparse category than adults did in Experiment 2A.

Method

Participants. Participants were sixty 4- and 5-year-olds ($M_{age} = 60$ months, $SD = 4.1$ months; 31 girls and 29 boys), with about equal numbers of children participating in each of the four conditions. An additional 28 children were tested and omitted from the sample because their performance in the catch trials did not meet the criterion (see *Materials, design, and procedure*).

Materials, design, and procedure. Materials and design were identical to those used in Experiment 2A, with the cover story identical to the one used in Experiment 1B. Pilot results with a separate group of 4- and 5-year-olds indicated that 24 learning trials would be sufficient for learning the dense category. Given that a lengthy procedure could be fatiguing for young children, the number of learning trials was limited to 24.

In the supervised learning condition, children were presented with the category-inclusion rule, identical to the one used for the adults in Experiment 2A. To make the mathematical relation among buttons, tails, and fingers more obvious, we arranged the relevant features depicted in the verbal description of the rule in two lines (one line of buttons and one line of tails and fingers together) with equal spacing between each of the elements. The line of buttons was visibly shorter than the line of tails and fingers. Similar to previous experiments, no pictures of creatures were presented in the supervised condition. Each rule (for the dense and sparse category) was repeated three times with three different depictions of the features.

The testing phase was identical to that in Experiment 2A, with the exception of having fewer testing trials (only 16 instead of 32) and fewer catch trials (six instead of eight). To be included in the study, children had to reject four of the catch trials.

Results and Discussion

Mean accuracy scores by category type and learning condition are presented in Figure 6B. A 2 (category type: dense vs. sparse) \times 2 (learning condition: unsupervised vs. supervised) between-subjects ANOVA rendered the predicted interaction significant, $F(1, 56) = 14.46$, $p < .001$. For the dense category, average accuracy scores were significantly higher in the unsupervised than the supervised condition, independent-sample $t(30) = 2.70$, $p < .01$, whereas for the sparse category, the scores were significantly higher in the supervised than the unsupervised condition, independent-sample $t(26) = 2.50$, $p < .02$. Furthermore, there was no evidence of learning the dense category in the supervised condition, and there was no evidence of learning the sparse category in the unsupervised condition (mean accuracy scores did not differ from zero; one-sample $t_s < 1.00$, $p_s > .34$). These results support the predictions revealing the same learning dissociation found for children in Experiment 1B and for adults in Experiments 1A and 2A: The unsupervised learning condition favored acquisition of the statistically dense category, and the supervised learning condition favored acquisition of the statistically sparse category.

Overall, Experiments 1–2 underscore the dissociation between statistically dense and sparse categories for both adults and children. It could be argued, however, that the reported findings reflect the difference in the overall number of relevant dimensions (or relations among dimensions) rather than statistical density, which is the ratio of relevant to irrelevant dimensions. Experiment 3 addressed this question by manipulating (a) the number of relevant dimensions, (b) the total number of dimensions, and (c) the statistical density of the category. If category learning is indeed a function of density, there should be a strong correlation between the statistical density of a category and its ease of acquisition.

Experiment 3

Method

Participants. Participants were 224 adults (127 women and 97 men), none of whom participated in any of the previous experiments. They were randomly assigned to one of the 15 conditions, with about equal numbers of participants in each condition. Twenty-one participants were tested and omitted from the sample because of their low performance on catch trials.

Materials and design. To eliminate potential differences in salience of dimensions, we created a new set of stimuli for which all varying dimensions pertained to the shading of shapes. The

stimuli were drawings of creatures that consisted of same-sized shapes (circles, triangles, and diamonds) of different colors and pattern (see Figure 7 for examples). To increase somewhat the within-category variability, the creatures had different kinds of line-drawn shoes.

Depending on condition, items could have 1, 3, 5, 7, or 10 shapes. Category membership was defined by color variation in all or just a subset of the shapes. Table 3 presents the resulting between-subject conditions as a function of (a) the total number of dimensions and (b) the number of relevant dimensions. The statistical density of the categories in each condition is given in the respective cell of Table 3. Note that the diagonal cells in Table 3 have categories with identical densities even though the categories differ in the total number of dimensions and the number of relevant dimensions. Acquisition of these categories should be equivalent if statistical density is the crucial factor.

Procedure. Participants had to distinguish between poisonous and harmless creatures. During training, participants were presented with 16 poisonous creatures (target category) and were instructed to pay particular attention to the color of the creature’s body parts. During the test phase, participants were presented with a new set of 16 items, half of which belonged to the target category and half of which belonged to the contrasting category (harmless creatures). To check overall alertness, we presented participants

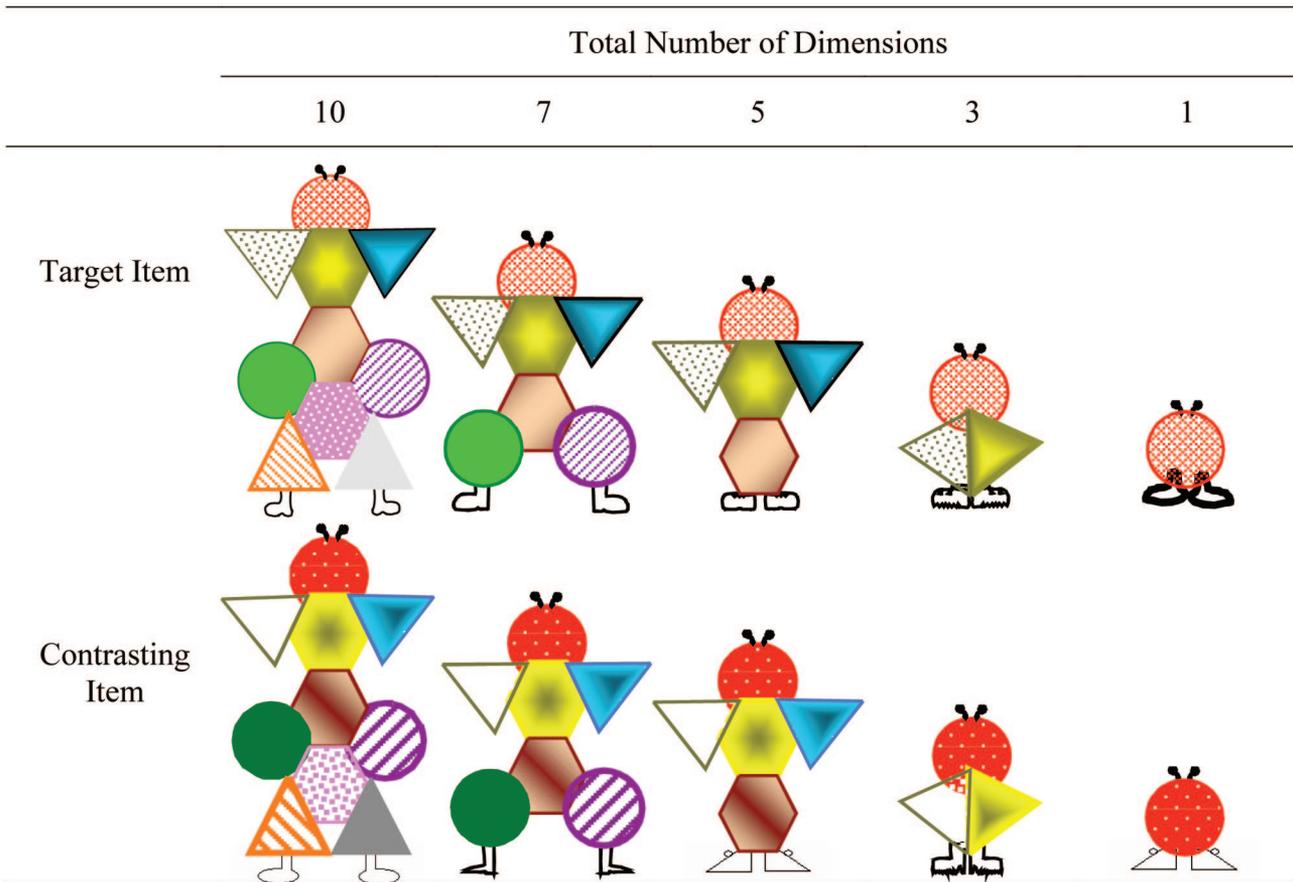


Figure 7. Examples of a stimuli used in Experiment 3 for the categories with density $D = 1.00$.

Table 3
Between-Subject Conditions of Experiment 3 Determined by the Number of Relevant Dimensions and the Total Number of Dimensions in a Category

Number of relevant dimensions	Total number of dimensions				
	10	7	5	3	1
10	1.00				
7	.70	1.00			
5	.50	.71	1.00		
3	.30	.43	.60	1.00	
1	.10	.14	.20	.30	1.00

Note. Values in cells represent the statistical density of a particular category.

with eight catch trials at the end of the testing phase. The catch items were creatures with a new number of body parts and new colors. To be included in the study, participants had to reject at least six out of the eight catch items.

Results and Discussion

Accuracy scores were calculated for each participant. The mean scores by condition are presented in Table 4. These scores were subjected to a regression analysis with the predictors density, total number of dimensions, and number of relevant dimensions. The analysis revealed a significant effect of the model, $F(3, 220) = 52.00, p < .001$, with density being the only significant predictor ($\beta = .62, t = 4.12, p < .0001$). Neither the total number of dimensions ($\beta = -.02, t = -1.54, p > .12$) nor the number of relevant dimensions ($\beta = .01, t < 1.00, p > .5$) was significant. These results undermine the possibility that participants' ability to learn the category was affected by extraneous variables that might sometimes correlate with statistical density. Not only was category density the best predictor of categorization accuracy, it also was a highly accurate predictor, with density accounting for 88% of variance in categorization accuracy (see Figure 8).

Experiment 4

Experiments 1–3 examined effects of statistical density on category learning. To summarize, we found that dense categories were ably learned without supervision, whereas sparse categories required explicit instruction. This was the case for both children and adults, and it applied to a variety of different categories:

Table 4
Mean Accuracy Scores as a Function of Condition Found in Experiment 3

Number of relevant dimensions	Total number of dimensions				
	10	7	5	3	1
10	.86				
7	.59	.94			
5	.31	.53	.81		
3	.31	.53	.52	.93	
1	.19	.22	.44	.41	.92

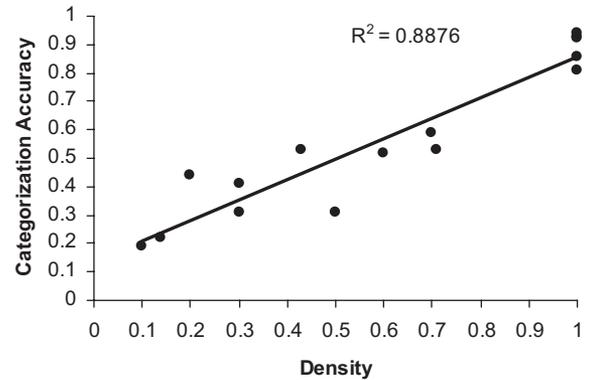


Figure 8. Categorization accuracy predicted by statistical density in Experiment 3.

categories that were defined by individual dimensions (Experiments 1 and 3) or categories defined by relations between dimensions (Experiment 2).

In Experiment 4, we examined the representation of dense and sparse categories in adults (Experiment 4A) and in children (Experiment 4B). Recall that several outcome contingencies are possible: Both dense and sparse categories could be represented in a rule-based manner, both dense and sparse categories could be represented in a similarity-based manner, or dense and sparse categories could be represented differently (e.g., the learner might form similarity-based representations for dense categories and rule-based representations for sparse categories).

Stimuli were colorful drawings of artificial creatures (similar to the ones used in Experiments 1 and 2) that varied in their appearance (hereafter, A) and in an arbitrary category-inclusion rule (hereafter, R) that could be manipulated independent of appearance. The resulting four types of test items were items for which appearance and rule matched those of the target category ($A_T R_T$), items for which appearance and rule matched those of the contrasting category ($A_C R_C$), and items for which either appearance alone or rule alone matched that of the target category ($A_T R_C$ or $A_C R_T$, respectively). The statistically dense category was defined by the overall appearance and the arbitrary rule, whereas the statistically sparse category was defined by the arbitrary rule alone.

As in the previous experiments, the procedure consisted of a learning phase and an immediate testing phase. However, in contrast to the previous experiments, both dense and sparse categories were acquired under the same learning regime, which included both explicit instruction and exposure to target items. Therefore, a difference in performance cannot be attributed to differences in the two learning procedures. The testing phase consisted of a surprise recognition task that included the target items (i.e., $A_T R_T$) and the three types of foils (i.e., $A_C R_C$, $A_T R_C$, and $A_C R_T$), none of which had been presented during training. If participants acquired the category at all, they should have correctly accepted $A_T R_T$ items and correctly rejected $A_C R_C$ items. Importantly, if they formed a representation based on the overall appearance, they should have false-alarmed on $A_T R_C$ items, but not on $A_C R_T$ items. If they formed a rule-based representation, they should have false-alarmed on $A_C R_T$ items, but not on $A_T R_C$ items. These predictions were tested with adults (Experiment 4A) and children (Experiment 4B).

Experiment 4A

Method

Participants. Participants were 26 adults (9 women and 17 men), none of whom participated in any of the previous experiments. They were randomly assigned to one of two conditions (learning the dense vs. sparse category), with about equal numbers of participants in each condition. Eight participants were tested and omitted from the sample because their performance on target items ($A_T R_T$) and contrasting items ($A_C R_C$) did not reach criterion (see *Procedure*).

Materials. Stimuli were artificial creatures similar to the ones used in Experiments 1–2. The varying features were the sizes of tail, wings, and fingers; the shadings of body, antenna, and buttons; and the numbers of fingers and buttons. The relation between the two latter features defined the arbitrary rule: Members of the target category had either many buttons and many fingers or few buttons and few fingers. All the other features constituted the appearance features. Members of the target category had a long tail, long wings, short fingers, dark antennas, a dark body, and light buttons (target appearance A_T), whereas members of the contrasting category had a short tail, short wings, long fingers, light antennas, a light body, and dark buttons (contrasting appearance A_C). Table 5 shows how appearance features and rule features were combined to create the four types of stimuli, and Figure 9 shows examples of each kind of stimulus. As in previous experiments, each appearance feature had two levels to increase variation between individual items.

Appearance features were probabilistic in each set of foils. For example, for $A_T R_T$ foils, only seven out of eight stimuli had a long tail, whereas the eighth stimulus had a short tail. This ensured that participants could not simply focus on a single feature to discriminate successfully between target and contrasting items. In contrast, the arbitrary rule was fully predictive. If participants simply focused on the most predictive feature, they should focus on the rule, and hence, they should false-alarm on $A_C R_T$, but not on $A_T R_C$ items.

Design and procedure. The experiment included two between-subjects conditions that differed in the density of the to-be-learned category. In the dense category, both appearance and rule were relevant for category membership, whereas, in the sparse category, only the inclusion rule was relevant for category membership. The procedure in these two conditions differed only in the explicit description of the category's inclusion rule presented to participants during training. In the dense-category condition, the description was "Most Ziblets have dark antennas, a dark body with light buttons, a long tail, and long yellow wings with short aqua fingers. Also, Ziblets with many aqua fingers on each yellow wing have many buttons, and Ziblets with few aqua fingers on each yellow wing have few buttons." In the sparse-category condition, the description was "Ziblets with many aqua fingers on each yellow wing have many buttons, and Ziblets with few aqua fingers on each yellow wing have few buttons." In addition to the verbal description of target items, participants were also presented with exemplars of the target category ($A_T R_T$) during training. These stimuli were identical for both conditions, so participants could focus on either appearance or rule features to distinguish between Ziblets and non-Ziblets.

After the training phase, participants were given a surprise recognition task: They were asked to distinguish between items that they had seen during the training and those that were new. The first set of test items consisted of eight $A_T R_T$ items and eight $A_C R_C$ items presented in a random order. The goal of this first set was to assess whether participants learned to discriminate target items from contrasting items. To be included in the study, participants had to perform correctly on at least 11 out of 16 trials (above chance, binomial test, $p > .06$). The second set of test items consisted of critical lures, with eight $A_T R_C$ items and eight $A_C R_T$ items. The goal of this set was to assess participants' representation of a category formed during category learning. Note that neither type of these stimuli was presented during training and that false alarms on any type of stimuli would give information about what was represented of the category of Ziblets.

Table 5
Examples of Stimuli Used in Experiment 4 Presented in Abstract Notation

Feature	$A_T R_T$		$A_T R_C$		$A_C R_T$		$A_C R_C$	
	Example 1	Example 2						
Appearance								
Size of tail	0	0	0	0	1	1	1	1
Size of wings	0	0	0	0	1	1	1	1
Size of fingers	1	1	1	1	1	0	0	0
Shading of body	0	0	0	0	1	1	1	1
Shading of antennas	0	0	0	0	1	1	1	1
Shading of buttons	1	1	1	1	1	0	0	0
Rule								
Number of fingers	0	1	0	1	0	1	0	1
Number of buttons	0	1	1	0	0	1	1	0

Note. There are two examples of each of the four item types. The numbers 0 and 1 refer to the values of the respective dimension (e.g., 0 = short tail, 1 = long tail). $A_T R_T$ items are referred to as target items, $A_C R_C$ items are referred to as contrasting items, and $A_T R_C$ and $A_C R_T$ items are critical lures. $A_T R_T$ = appearance and rule match those of target category; $A_T R_C$ = appearance alone matches that of target category; $A_C R_T$ = rule alone matches that of target category; $A_C R_C$ = appearance and rule match those of contrasting category.

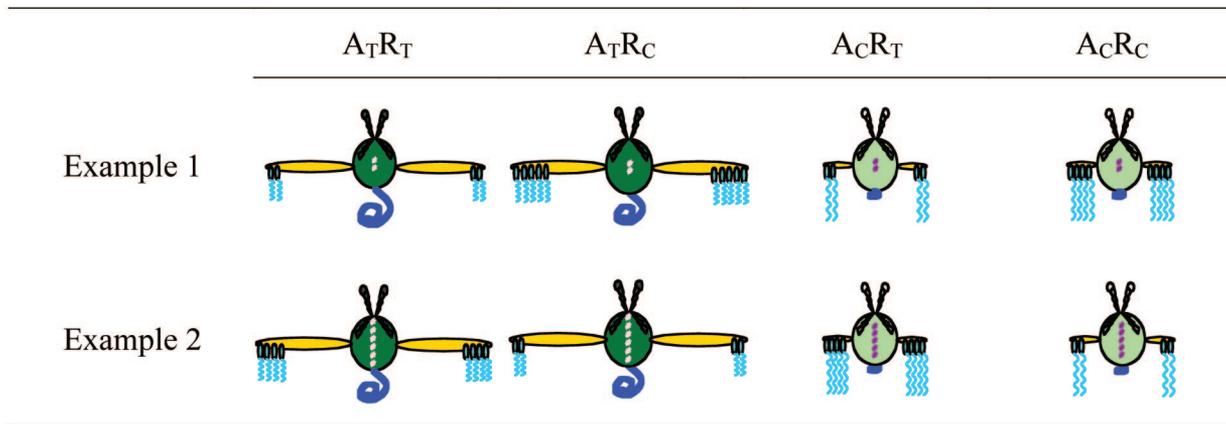


Figure 9. Examples of stimuli used in Experiment 4. $A_T R_T$ = appearance and rule match those of target category; $A_T R_C$ = appearance alone matches that of target category; $A_C R_T$ = rule alone matches that of target category; $A_C R_C$ = appearance and rule match those of contrasting category.

Results and Discussion

Across conditions, $A_T R_T$ items were accurately accepted at $M = 0.90$, and $A_C R_C$ items were accurately rejected at $M = 0.91$ with no significant difference between the conditions ($ps > .22$). We therefore focus on participants' performance on critical lures. Figure 10A presents the mean proportion of "yes" responses (i.e., false alarms) broken down by the foil type ($A_T R_C$ vs. $A_C R_T$) and condition (dense vs. sparse). As can be seen in the figure, in the dense-category condition, participants were more likely to false-alarm on items that were similar to training items, whereas, in the sparse-category condition, they were more likely to false-alarm on items that had the same rule as the training items.

A 2 (category type: dense vs. sparse) \times 2 (foil type: $A_T R_C$ vs. $A_C R_T$) mixed-design ANOVA revealed a significant interaction, $F(1, 24) = 18.40$, $p < .001$. In particular, participants in the dense-category condition were more likely to false-alarm on items that shared the appearance of studied items than on items that shared the category-inclusion rule, paired-sample $t(11) = 2.37$, $p < .05$. The opposite was the case for the sparse-category condition, with participants being more likely to false-alarm on items that shared the rule of the learned category than on items that shared its appearance, paired-sample $t(13) = 3.97$, $p < .01$. Recall that both pieces of information—appearance and arbitrary rule—were available to participants during training, with the inclusion rule of the category being the only difference between the conditions. Thus, when presented with the inclusion rule of a sparse category, participants tended to filter out appearance information presented during training and formed a rule-based category representation. In contrast, when presented with the inclusion rule of a dense category, participants were more likely to focus on appearance, forming a similarity-based representation. These findings indicate that the distinction between dense and sparse categories is not limited to learning but is also evident in how participants represent categories. The goal of Experiment 4B was to examine whether the same holds for young children.

Experiment 4B

Method

Participants. Participants were twenty-two 4- and 5-year-olds ($M_{age} = 59$ months, $SD = 3.5$ months; 12 girls and 10 boys), randomly assigned to one of the two conditions. Four additional children were tested and omitted from the sample because their accuracy on $A_T R_T$ and $A_C R_C$ foils was below 0.5.

Materials, design, and procedure. Materials and design were identical to those used in Experiment 4A, and the cover story presented to children was identical to that used in Experiment 1B and 2B.

Results and Discussion

Across conditions, $A_T R_T$ items were accurately accepted ($M = 0.78$), and $A_C R_C$ items were accurately rejected ($M = 0.91$), with no significant difference between the conditions ($ps > .46$). Figure 10B presents children's mean proportion of "yes" responses (i.e., false alarms) on critical lures ($A_T R_C$ vs. $A_C R_T$) as a function of condition (dense vs. sparse). A 2 (category type: dense vs. sparse) \times 2 (foil type: $A_T R_C$ vs. $A_C R_T$) mixed-design ANOVA revealed only a significant effect of foil type, $F(1, 20) = 103.60$, $p < .001$, with children being more likely to false-alarm on $A_T R_C$ foils ($M = 0.66$) than on $A_C R_T$ foils ($M = 0.29$). This finding indicates that children represented categories by appearance and not by the rule whether the category was dense or sparse.

One could argue, however, that children represented the appearance of the sparse category because they failed to learn the inclusion rule of the sparse category. In particular, children tested in the sparse-category condition might have had difficulty with the verbal description of the arbitrary rule R_T and therefore might have focused only the appearance of the items presented to them during training. To rule out this possibility and to examine children's ability to learn the category, we tested a new group of children ($n = 33$, mean age = 59.7 months; 16 girls and 17 boys). The same materials and procedure were used, with the only difference being

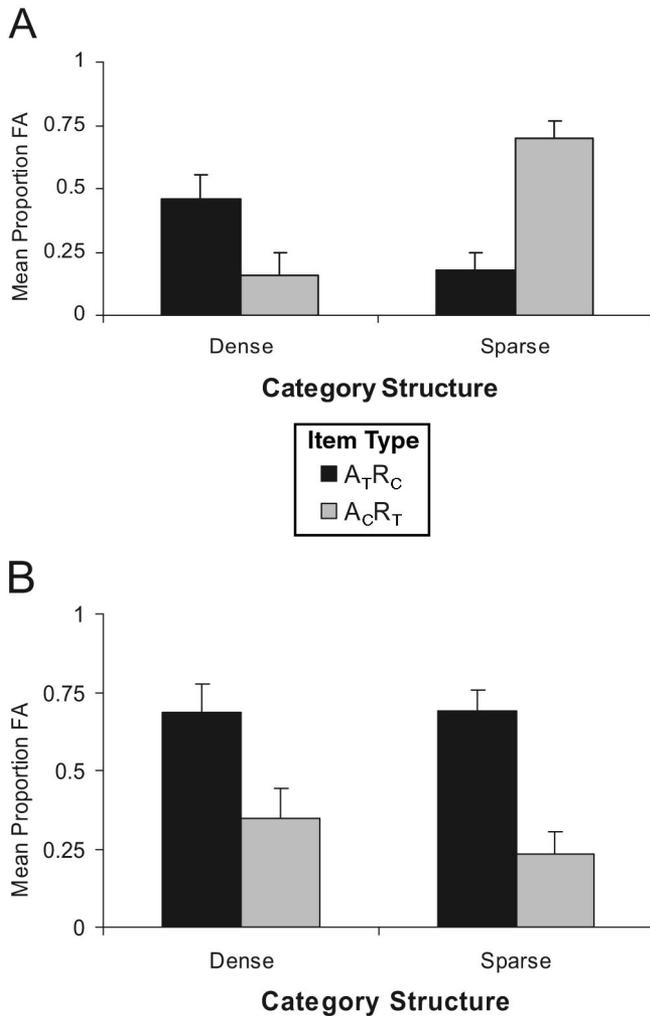


Figure 10. Mean accuracy scores by item type and category structure in Experiments 4A and 4B. Error bars represent standard errors of the mean. FA = false alarms; A_TR_C = appearance alone matches that of target category; A_CR_T = rule alone matches that of target category.

that children were asked to perform a categorization task (i.e., distinguish between Ziblets and Flurps), not a recognition task (i.e., distinguish between old and new items). Learning of the dense category would yield acceptance of A_TR_C, but not A_CR_T, items as Ziblets, whereas learning of the sparse category would yield acceptance of A_CR_T, but not A_TR_C, items as Ziblets. This interaction between category type (dense vs. sparse) and foil type (A_TR_C vs. A_CR_T) was indeed found, $F(1, 31) = 13.20, p < .01$, with children in the dense-category condition being more likely to categorize A_TR_C items than A_CR_T items as Ziblets ($M = 0.64$ vs. $M = 0.22$), paired-sample $t(10) = 3.10, p < .05$, and children in the sparse-category condition being more likely to categorize A_CR_T items than A_TR_C items as Ziblets ($M = 0.71$ vs. $M = 0.22$), paired-sample $t(21) = 16.50, p < .01$. This finding suggests that children can use the rule of the sparse category when asked to categorize items even though they represent the category in terms of its appearance.

Taken together, results of Experiments 4 indicate that adults and children differ in how statistical density affects category represen-

tation. Adults represent sparse categories in a rule-based manner, and they represent dense categories in a similarity-based manner. In contrast, young children represent both dense and sparse categories in a similarity-based manner. These findings indicate that the similarity-based representation is a developmental default, whereas rule-based representations are a product of learning and development. Overall results of Experiments 4A and 4B present interesting challenges for theories of categorization to explain the developmental transition, and they may have implications for research on memory and memory distortions.

General Discussion

Results of the reported experiments point to important dissociation between dense and sparse categories with respect both to category learning and to category representation. In Experiment 1, the dissociation was found for categories with linearly separable category structures in children and adults: Although supervision did not facilitate learning of dense categories, supervision markedly increased learning of sparse categories. Experiment 2 corroborated these findings with nonlinearly separable categories again in children and adults. Furthermore, Experiment 3 found that density predicts unsupervised category learning markedly better than other predictors, such as the number of relevant dimensions or the total number of dimensions.

In Experiment 4, the dissociation between dense and sparse categories was tested in terms of category representation after dense and sparse categories were learned under comparable learning regimes. Adults were more likely to represent sparse categories in a rule-based manner, whereas they were more likely to represent dense categories in a similarity-based (or appearance-based) manner. Specifically, when categories were dense, adults were likely to false-alarm on items that shared the target appearance rather than on items that shared the arbitrary rule of the target, whereas the reverse was the case for sparse categories. This pattern of responses was not the case for young children: Unlike adults, children represented dense and sparse categories on the basis of the category's similarity. Specifically, they were more likely to false-alarm on items that shared the target appearance than on items that shared the arbitrary rule of the target regardless of category structure. This finding indicates that rule-based representations of sparse categories found in adults are not a default but are rather a product of development.

Statistical Density: A Solution to the Learning Paradox?

Results of current research offer a potential solution to the paradox of category learning described in the introductory section, above. Recall that although some categories are easily learned without supervision even by young infants, other categories require supervised effortful learning even in adults. Research presented here suggests that category density is a candidate solution to the paradox: Whereas dense categories can be easily acquired without supervision, learning of sparse categories requires supervision.

Results of Experiment 1C support our hypothesis that learning of dense categories puts small demands on selective attention (because most of the information is category relevant) and thus can be learned without supervision. In particular, when learning dense categories in Experiment 1C, learners attended to the entire pattern rather than to a single dimension. In contrast, learning statistically

sparse categories requires the learner to ignore a large amount of category-irrelevant information while focusing on category-relevant information. The greater the proportion of the to-be-ignored information (i.e., the sparser the category), the more difficult it is to figure out what should be ignored and, as a result, the more difficult it is to learn the sparse category without supervision. Experiments 1B, 2A, and 2B support this contention, indicating that children and adults could not learn the respective sparse categories without supervision. Yet even children could learn the sparse categories when learning was supervised. Taken together, these findings support the prediction that category structure interacts with the learning regime.

Study of categorization has primarily relied on supervised category learning (Love, 2002). Therefore, theories formed on the basis of evidence from supervised learning may tend to overestimate the difficulty of learning of denser categories, especially when these categories are probabilistic. Love (2002) compared acquisition of Shepard et al.'s (1961) category types across different learning modes and found that the categories that generated substantial difficulty under supervised learning were more easily acquired under unsupervised, incidental learning. Similarly, theories relying primarily on evidence from unsupervised learning (e.g., Whitman & Garner, 1962) may overestimate the difficulty of learning sparse categories: Acquisition of some sparse categories may be exceedingly difficult (if not impossible) under an unsupervised learning regime, whereas these same categories can be readily acquired under supervised learning.

The concept of statistical density also brings a principled solution to the so-called feature-selection problem. This problem is based on the idea that many categories have multiple features and feature relations, with many being irrelevant for category membership. For example, even though many refrigerators are white, membership in the category *refrigerator* is not affected by color. It has been argued that because category learning requires attention to relevant features and relations, the learner has to know in advance what the relevant features are. This means that category learning requires some form of top-down knowledge (see Murphy & Medin, 1985, for arguments). The current research suggests that although there might be a problem of feature selection when categories are sparse, this is not the case for dense categories. The highly intercorrelated structure of dense categories renders top-down knowledge unnecessary: Because dense categories have many redundant features, there is no need to selectively attend to a particular feature. This could enable even young infants to effortlessly acquire dense categories. It can also be successfully exploited by neural network models of conceptual development (e.g., Rogers & McClelland, 2004) because even simple networks can acquire dense categories without top-down knowledge of which features are important.

Density and Representation of Categories

It has been long believed that most (if not all) categories are represented in a similar way (for a review, see Murphy, 2002; E. E. Smith & Medin, 1981). For example, according to the classical view, category learning is viewed as a process of discovery of the necessary and sufficient features, with necessary and sufficient features being central for category learning and representation (Bruner et al., 1956; Vygotsky, 1986). Alternatively, it has been argued that categorization is grounded in perceptual and attentional mechanisms capable of detecting similarities in the input

(e.g., Goldstone, 1994; Hampton, 1998; Kruschke, 1992; Rogers & McClelland, 2004; Sloutsky & Fisher, 2004; for reviews, see Hahn & Ramscar, 2001; Rips, 2002; Sloman & Rips, 1998). According to that account, categories are represented as a central tendency (Minda & Smith, 2001; Posner & Keele, 1968; Rosch, 1978), as probabilistic rules (Ashby, 1989), or as sets of encountered exemplars (Hintzman, 1986; Kruschke, 1992; Medin & Schaeffer, 1978; Nosofsky, 1986, 1992).

However, there is a growing body of evidence that not all categories are the same, and it is possible that each of these theories is partially correct as they accurately describe representation of different types of categories. This idea has given rise to a variety of multisystem/multiprocess models that assume a combination of unitary processes, some of which are rule-based and some of which are similarity-based processes (e.g., Ashby et al., 1998; Erickson & Kruschke, 1998; Nosofsky et al., 1994; E. E. Smith & Sloman, 1994).

Our results indicate that in adults, dense categories are represented in a similarity-based manner and sparse categories are represented in a rule-based manner. These findings are consistent with the multisystem/multiprocess models, although it is conceivable that a single-process model would also be able to account for these findings. The developmental differences in category representation found here suggest that similarity-based representation is a developmental default, whereas rule-based representations are a product of development. These results support previous developmental findings (e.g., Sloutsky & Fisher, 2004; see also Sloutsky, 2003, and Sloutsky & Fisher, 2005, for discussions) while posing an interesting challenge for theories of categorization: Categorization theories have to account for the developmental difference in how sparse categories are represented, and they have to address what underlying processes could explain the developmental transition.

Statistical Density: Just Another Distinction?

As noted in the introductory section, above, there have been a number of proposals casting doubt on the assumption that all or most categories are the same (see Gentner & Kurtz, 2005; Goldstone, 1996; E. E. Smith, Patalano, & Jonides, 1998, for reviews). How then does the concept of statistical density add to these already existing distinctions?

It is worth noting that statistical density maps well (albeit not perfectly) onto many of the previously proposed distinctions. In particular, it maps well onto Rosch's taxonomic distinction (Rosch & Mervis, 1975), given that members of basic-level categories (e.g., *bird*) have greater statistical density than members of superordinate-level categories (e.g., *animal*). Furthermore, it maps well onto the concrete-versus-abstract distinction given that concrete categories (e.g., *car*) are likely to embody entities with a higher number of overlapping perceptual features than abstract categories (e.g., *truth*). Statistical density also maps well onto the entity-versus-relational-category distinction and hence onto the noun-versus-verb distinction (e.g., Gentner, 1981). Entity categories (e.g., *car*) are likely to be dense, and relational categories (e.g., *driving*) are likely to be sparse (see Gentner & Kurtz, 2005, for related arguments). Finally, it maps onto the distinction between natural and nominal kinds, with natural kinds (e.g., species of animals) being statistically denser than nominal kinds (e.g., scientific, mathematical, and legal concepts).

It is possible that the concept of statistical density is a more general distinction, one that lies at the basis of other distinctions. One important advantage of statistical density is that, unlike some of the other proposed measures, statistical density can be measured independently (in principle) rather than be inferred from participants' patterns of response. It furthermore provides a continuous measure rather than a mere dichotomy. As such, it makes it possible to capture the graded nature of differences between categories.

Category Structure and Learning

There are several theoretical models of statistical structure that predict effects of statistical structure on category learning. In particular, Garner (1962) argued that learning of a set of stimuli is a function of feature redundancy within this set. Similarly, Trabasso and Bower (1968) argued that the proportion of category-relevant to category-irrelevant information predicts the efficiency of category learning. Homa et al. (1979) demonstrated that category learning is affected by the structural ratio of a category or category coherence (see also J. D. Smith & Minda, 2000), a measure that is based on a summed within- and between-category distances of items in the psychological space. Corter and Gluck (1992) introduced the idea of category utility to capture differences in induction efficiency, with category utility being a function of a feature's frequency in a category (or category validity) and the base rate of the feature in predicting feature induction.

Although some of these models predict effects of category structure on category learning, none of the models predicts the interaction of category structure with the learning regime. In particular, some models (e.g., Garner, 1962; Homa et al., 1979) can predict better learning of more redundant categories, but they cannot account for able learning of exceedingly sparse categories. This is because their models of category structure do not include selective attention. As a result, these models cannot account for our findings that even exceedingly sparse categories (such as those used in Experiment 2) can be learned under supervision.

Other models, such as that of Trabasso and Bower (1968), can account for successful supervised learning of dense categories while having difficulty accounting for unsupervised learning of these categories. This is because these models assume that participants sample and test hypotheses about category-relevant dimensions, with both sampling and testing requiring feedback. However, results of Experiment 1C cast doubt on this assumption: Participants did not seem to sample individual dimensions when learning dense categories. More importantly, results of Experiments 1–3 reveal efficient category learning under the feedback-free unsupervised learning regime. In sum, whereas previous models of category structure can predict only a main effect of category structure on learning, the current work suggests an interaction between category structure and learning regime.

Questions for Future Research

Despite these advances, several questions remain. In particular, how much statistical density is needed for the category to be learned without supervision (in trivial time)? Does this amount change in the course of development? How do different types of supervised learning affect acquisition of sparse categories? Recall that the current research examined only one type of supervision, namely, explicit instruction. It

would be important to examine how more subtle types of supervision (e.g., negative evidence, corrective feedback, or guided comparison) affect acquisition of dense and sparse categories. These questions have to be addressed in future studies.

Conclusions

The results of the current research reveal several important regularities about category learning and representation in children and adults. First, sparse, but not dense, categories require supervised learning. Second, later in development, sparse categories are represented in a rule-based manner, whereas dense categories are represented in a similarity-based manner. Finally, early in development, both dense and sparse categories are represented in a similarity-based manner, thus suggesting that similarity-based category representation is a developmental default. Taken together, these findings provide important insights about the role of category density in category learning and category representation across development.

References

- Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, *120*, 3–19.
- Ashby, F. G. (1989). Stochastic general recognition theory. In D. Vickers & P. L. Smith (Eds.), *Human information processing: Measures, mechanisms, and models* (pp. 435–457). Amsterdam: Elsevier.
- 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.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–609.
- Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 458–475.
- Bomba, P. C., & Siqueland, E. R. (1983). The nature and structure of infant form categories. *Journal of Experimental Child Psychology*, *35*, 294–328.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York: Wiley.
- Corter, J., & Gluck, M. (1992). Explaining basic categories: Feature predictability and information. *Psychological Bulletin*, *111*, 291–303.
- Dempster, F. N., & Corkill, A. J. (1999). Interference and inhibition in cognition and behavior: Unifying themes for educational psychology. *Educational Psychology Review*, *11*, 1–88.
- Eimas, P. D., & Quinn, P. C. (1994). Studies on the formation of perceptually based basic-level categories in young infants. *Child Development*, *65*, 903–917.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.
- French, R. M., Mareschal, D., Mermillod, M., & Quinn, P. C. (2004). The role of bottom-up processing in perceptual categorization by 3- to 4-month-old infants: Simulations and data. *Journal of Experimental Psychology: General*, *133*, 382–397.
- Garner, W. R. (1962). *Uncertainty and structure as psychological concepts*. New York: Wiley.
- Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, *4*, 161–178.
- Gentner, D., & Kurtz, K. J. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 151–175). Washington, DC: American Psychological Association.
- Goldstone, R. L. (1994). Similarity, interactive activation, and mapping.

- Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 3–28.
- Goldstone, R. L. (1996). Isolated and interrelated concepts. *Memory & Cognition*, 24, 608–628.
- Hahn, U., & Ramscar, M. (Eds.). (2001). *Similarity and categorization*. London: Oxford University Press.
- Hampton, J. A. (1998). Similarity-based categorization and fuzziness of natural categories. *Cognition*, 65, 137–165.
- Hintzman, D. L. (1986). “Schema abstraction” in a multiple-trace memory model. *Psychological Review*, 93, 411–428.
- Homa, D., Rhoades, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 11–23.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science*, 6, 449–476.
- Kripke, S. (1972). Naming and necessity. In D. Davidson & G. Harman (Eds.), *Semantics of natural language* (pp. 253–355). Dordrecht, Holland: Reider.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, 9, 829–835.
- Mareschal, D., & Quinn, P. C. (2001). Categorization in infancy. *Trends in Cognitive Sciences*, 5, 443–450.
- Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? *Annual Review of Psychology*, 51, 121–147.
- Medin, D. L., & Schaeffer, 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.
- 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. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.
- Napolitano, A. C., & Sloutsky, V. M. (2004). Is a picture worth a thousand words? The flexible nature of modality dominance in young children. *Child Development*, 75, 1850–1870.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nosofsky, R. M. (1992). Exemplars, prototypes, and similarity rules. In A. Healy, S. Kosslyn, & R. Shiffrin (Eds.), *From learning theory to connectionist theory: Essays in honor of W. K. Estes* (Vol. 1, pp. 149–168). Hillsdale, NJ: Erlbaum.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353–363.
- Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, 22, 463–475.
- Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 88–94.
- Rips, L. J. (2002). Reasoning. In H. Pashler & D. Medin (Eds.), *Steven’s handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (3rd ed., pp. 363–411). New York: Wiley.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Rouder, J. N., & Ratcliff, R. (2004). Comparing categorization models. *Journal of Experimental Psychology: General*, 133, 63–82.
- Shannon, C. E., & Weaver, W. (1948). *The mathematical theory of communication*. Chicago: University of Illinois Press.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, 75, 517.
- Sloman, S. A., & Rips, L. J. (1998). Similarity as an explanatory construct. *Cognition*, 65, 87–101.
- Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*, 7, 246–251.
- Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, 133, 166–188.
- Sloutsky, V. M., & Fisher, A. V. (2005). Similarity, induction, naming, and categorization (SINC): Generalization or verbal inductive reasoning? Response to Heit and Hayes (2005). *Journal of Experimental Psychology: General*, 134, 606–611.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition*, 65, 167–196.
- Smith, E. E., & Sloman, S. A. (1994). Similarity- versus rule-based categorization. *Memory & Cognition*, 22, 377–386.
- 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.
- Smith, L. B. (1989). A model of perceptual classification in children and adults from global similarities to kinds of similarities: The construction of dimensions in development. *Psychological Review*, 96, 125–144.
- Trabasso, T. R., & Bower, G. H. (1968). *Attention in learning: Theory and research*. New York: Wiley.
- Vygotsky, L. S. (1986). *Thought and language*. Cambridge, MA: MIT Press.
- Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, 8, 168–176.
- Whitman, J. R., & Garner, W. R. (1962). Free recall learning of visual figures as a function of form of internal structure. *Journal of Experimental Psychology*, 64, 558–564.
- Young, M. E., & Wasserman, E. A. (2001). Entropy and variability discrimination. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 278–293.
- Younger, B. A. (1993). Understanding category members as “the same sort of thing”: Explicit categorization in ten-month infants. *Child Development*, 64, 309–320.
- Younger, B. A., & Cohen, L. B. (1986). Developmental change in infants’ perception of correlations among attributes. *Child Development*, 57, 803–815.
- Zelazo, P. D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development*, 11, 37–63.

(Appendixes follow)

Appendix A

Estimating Weights of Individual Dimensions and Binary Relations Between Dimensions

To measure the degree to which isolated dimensions differ from dimension relations in terms of their salience, we asked adults ($n = 16$) to report the rule of various categories that were created either on the basis of a single dimension or on the basis of a dyadic relation. Target items and contrasting items that differed in the binary dimensions of color, shape, and size were created. Participants had full view of target items and contrasting items when they

had to generate the rule. The mean proportion of correct responses was 0.81 for categories created on the basis of a single dimension and 0.42 for categories created on the basis of a dyadic relation. These findings suggest the likelihood of detecting a relation is about a half of that of a dimension, at least under the favorable conditions under which both target and contrasting items are presented simultaneously.

Appendix B

Measuring Density for Categories Used in Experiments 1–3

Density across the reported experiments	H^{dim}	H^{rel}	$H^{dim} + H^{rel}$
Experiment 1: Dense category ($M = 6, O = 15$)			
Within-category	1 (6*0)	0.5 (15*0)	0
Between-category	1 (6*1)	0.5 (15*2)	21.0
$D = 1 - (0/21) = 1.00$			
Experiment 1: Sparse category ($M = 6, O = 15$)			
Within-category	1 (5*1 + 1*0)	0.5 (10*2 + 5*1)	17.5
Between-category	1 (6*1)	0.5 (15*2)	21.0
$D = 1 - (17.5/21) = 0.17$			
Experiment 2: Dense category ($M = 8, O = 28$)			
Within-category	1 (8*1)	0.5 (28*1)	22.0
Between-category	1 (8*1)	0.5 (28*2)	36.0
$D = 1 - (22/36) = 0.39$			
Experiment 2: Sparse category ^a			
Within-category	1 (8*1)	0.5 (27*2 + 1*0)	35.0
Between-category	1 (8*1)	0.5 (28*2)	36.0
$D = 1 - (35/36) = 0.03$			
Experiment 3: Category with five relevant dimensions ($M = 10, O = 45$) ^b			
Within-category	1 (5*1 + 5*0)	0.5 (10*2 + 25*1 + 10*0)	27.5
Between-category	1 (10*1)	0.5 (45*2)	55.0
$D = 1 - (27.5/55) = 0.50$			

Note. For simplicity, we assume that the attentional weight is 1.0 for a dimension and 0.5 for a relation. These weights are shown in front of the parentheses. The values in parentheses represent $-\sum (p_j \log_2 p_j)$ for dimensions and $-\sum p_{mm} \log_2 p_{mm}$ for relations, whereas values in italics represent $-p_i \log_2 p_j$ for dimensions and $-p_{mm} \log_2 p_{mm}$ for relations. M = total number of varying dimensions; O = total number of dyadic relations between two dimensions; H^{dim} and H^{rel} = entropy due to dimensions and due to relations, respectively; D = statistical density.

^a Calculations pertain to the sparse category used in the follow-up experiment of Experiment 2A.

^b Variation in creatures' shoes is ignored because of zero entropy for this dimension.

Received August 5, 2004
 Revision received February 8, 2007
 Accepted February 12, 2007 ■