Domain-specific knowledge in simple categorization tasks

DEBORAH KELEMEN and PAUL BLOOM
University of Arizona, Tucson, Arizona

Many contemporary theories of learning and memory adopt the empiricist premise that concepts are structured according to perceptual similarity. Developmental differences in categorization tasks are thereby interpreted as the result of qualitative shifts in the capacity to attend to specific perceptual dimensions. An alternative theory is that domain-specific knowledge underlies categorization, and that even performance on simple categorization tasks is influenced by such knowledge. To test this hypothesis, adults were asked to categorize colored circles, which were described as either natural kinds or artifacts. In Study 1, the subjects were shown the actual circles; in Study 2, they were given descriptions. Adults categorized the circles differently as a function of how they were described and were influenced on subsequent choices by the demand to create "cohesive" categories. These results suggest that the developmental shifts may be due to differences in domain-specific knowledge about the nature of categorization tasks, not due to global cognitive changes. This proposal is supported by evidence from previous studies of adult categorization and children's acquisition of word meaning.

Many contemporary theories of mind posit that concepts are abstractions from perceptual experience. Porcupines, for instance, fall into a different psychological category than cacti because they look different from cacti; they differ with regard to perceptual features or dimensions. Within philosophy, this view has been defended by empiricists such as Quine (1960); within psychology, the premise that concepts emerge out of some perceptual similarity space forms the basis of many classical and connectionist models of memory and learning (e.g., McClelland & Rumelhart, 1986).

One alternative to this is a "rationalist" perspective, wherein concepts are thought to emerge from domain-specific systems of knowledge, sometimes described as "stances" or "naive theories" (see, e.g., Bloom, in press; Carey, 1985; see also Murphy, 1993, for a review). Under this view, perceptual similarity is assumed to play a role in categorization only to the extent that things which are perceptually similar tend to share other, more important properties. As such, similarity of appearance is a *cue* to sameness of category, but it is not in itself criterial. Instead, notions about properties of entities such as their intended function (for artifacts) or their internal structure (for natural kinds) lie at the core of categorization. Porcupines and cacti fall into psychologically distinct categories then, in part because of our intuitions about deeper properties of these entities. In support of this, young children are quite willing to accept that something can be a porcupine even if it looks like a cactus, as long as they have reason to believe that its cactus-like appearance is the result of a set of superficial transformations (Keil, 1989; see also Gelman & Markman, 1987).

A phenomenon that is often interpreted as lending support to the empiricist position is the existence of developmental changes in performance on perceptual classification tasks. In an important series of studies, Smith (1989) presented subjects of different ages with circles of different sizes and colors that had been structured as in Figure 1. Stimuli A and B are identical on dimension X but differ along dimension Y, while stimulus C is slightly different from A on both dimensions. Importantly, A and C are more "overall similar" (OS) to one another than are A and B. In half the trials, the dimension of identity for A and B was color (ID[C]); in half, it was size (ID[S]). Subjects were asked to "Make groups, [to] put together the ones that go together."

While adults tended to base their judgments on identity (selecting A and B), Smith (1989) found, using a range of different stimuli items and configurations, that young children attended to overall similarity (selecting A and C). Smith's explanation for the developmental differences had two components. First, she argued that children's holistic strategy is the result of their difficulties in selectively attending to specific perceptual dimensions. This conclusion is further supported by evidence from a variety of sources, including findings that children have difficulties in attending to relevant dimensions in discrimination tasks and in identifying the dimension of change after watching an object undergo a
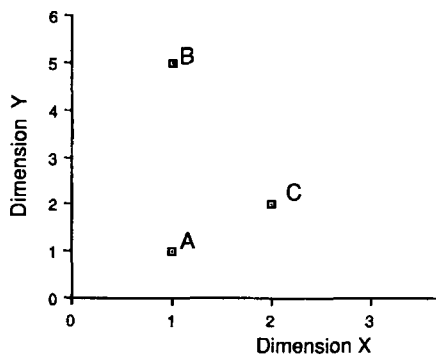
Figure 1. Structure of "standard triad" used by Smith (1989).

transformation (see Kemler, 1983, for review). Second, Smith argued that adults, but not children, are strongly motivated to classify by identity because only adults construe identity as a special relationship from the standpoint of inference. Evans and Smith (1988) argued that identity is special in that it licenses certain inferences that similarity does not. Transitivity is one example: If A is identical to B and B is identical to C, then A is identical to C; but if A is similar to B and B is similar to C, it need not be the case that A is similar to C.

An alternative explanation is that the difference in development is not the result of shifts (either global or domain by domain) in the capacity to attend to specific dimensions or to appreciate the value of identity. It is instead the result of differences in certain aspects of domain-specific knowledge that underlie the understanding of concepts. Such knowledge clearly has a role to play in more contentful categorization situations. For one thing, it can determine which dimensions are relevant in a classification. When one is given a new perfume and asked to find another of the same kind, for instance, knowledge about the function of perfume tells one that smell is relevant, but that the color or size of the bottle is not (see Pazzani, 1991). Furthermore, background knowledge can determine whether or not identity along a given perceptual dimension is especially significant. If there is little or no variation in the size of an object within a given category, as with some artifacts that are intentionally made with attention to size, then identity could be crucial—thus the function and manufacture of coins entails that if two coins are even slightly different in diameter it is a good bet that they belong to different types (see also Rips, 1989). For categories such as biological kinds where some arbitrary variation along such dimensions is to be expected, identity becomes less important.

One important finding along these lines is provided by Nisbett, Krantz, Jepson, and Kunda (1983). They asked subjects to consider a single exemplar of a category (birds or chemical elements) and estimate the likelihood that an identical trait (color or conductivity) would be present in other category members. Subjects were far more conservative in their estimates when con-

sidering the biological kind (where variance was expected with respect to color) than when considering the chemical element (which was expected to be invariant with respect to conductivity). To anticipate a point that will be made below, note that one would *not* expect to find accidental variation in the size and color of stimuli that have been intentionally constructed by an experimental psychologist.

To examine the role of domain-specific knowledge in classification, we tested adults on simple categorization tasks modeled on those used by Smith (1989). The goal of these experiments was to explore the nature of the *adult* categorization process, because an understanding of mature competence in a given domain is an essential aspect of any theory of developmental changes within that domain. The rationalist perspective on adult categorization makes two predictions about adult performance on these tasks.

Prediction 1. As a function of how a set of colored circles is described, subjects will focus on different dimensions in a classification task. In particular, when the circles are described as microscopic animals, subjects will focus on the dimension of color because of their belief that animals from the same group should share the same approximate pigmentation. Size is less relevant, because it shows more variation, changing over the animal's lifetime. In contrast, for circles described as tiny machines, size will be considered the more valid basis for classification, because size, but not color, is relevant to machine function. A model of perceptual classification such as Smith's (1989) predicts that subjects should either attend to both dimensions equally or consistently favor one dimension over the other, regardless of the background description.

Prediction 2. Once subjects have selected a particular dimension as relevant for a first choice, they will continue to attend to this dimension for future choices when forming the same category. For instance, once subjects classify two items together because they have the same color, they will prefer to classify a third item as belonging to this group if it is also similar in color—even in the presence of another stimulus that is identical in size. In order to create a coherent category, then, subjects could end up disregarding identity and favoring nonidentity relations along an already-chosen perceptual dimension. A natural extension of the Evans and Smith (1988) proposal makes the alternative prediction; adults will tend to consistently value identity over nonidentity in classification because of the inferential power of the identity relationship.

STUDY 1

Method

Subjects. Thirty-two undergraduates from the University of Arizona volunteered for this study.

Materials and Procedure. Classification sets were constructed according to the schemes in Figure 2. The stimuli were circles of Coloraid paper on white cards. There was a target item, an item identical in color to the target but very different in size (ID[C]), an item identi-
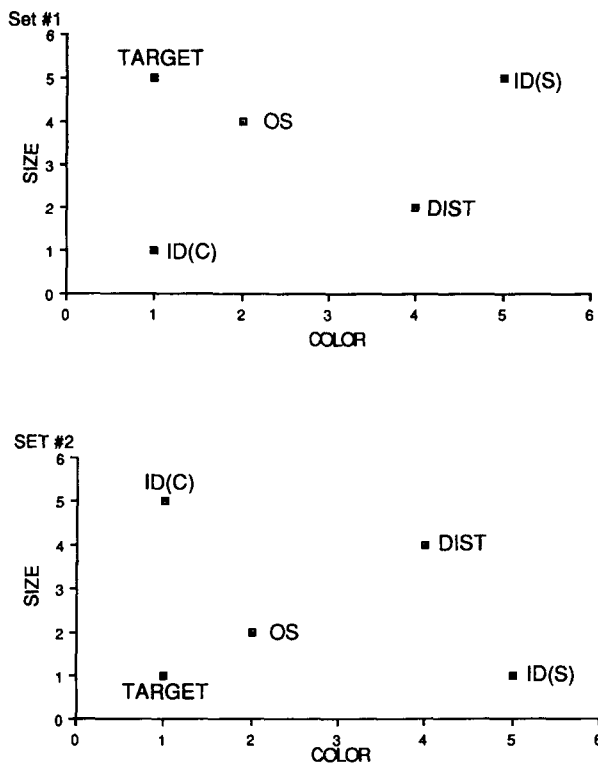
Figure 2. Structure of stimuli used in Study 1 and Study 2.

cal in size to the target but very different in color (ID[S]), an item that was similar in both color and size, and—as determined in pretesting—overall similar to the target relative to the other items (OS), and a distractor that was very different in both color and size from the target (DIST) (see Kelemen, 1992, for discussion of the pretesting).

Each subject was given four sets of stimuli, two where the target was a small light-colored circle (diameter of 3.25 cm), and two where the target was a large dark-colored circle (diameter of 6.75 cm). Half the sets consisted of green circles of different shades; half consisted of blue circles.

Subjects were presented with four trials, which alternated between blue and green classification sets. Half the subjects heard the stimuli described as "microscopic animals"; the other half heard them described as "tiny machines." The machines/animals in both blue circle trials were given one name (intended to be interpreted as a superordinate); those in the green circle trials were given a different name. The subjects were asked to imagine that they had taken a summer job helping scientists (for the animals) or engineers (for the machines) by looking at and classifying these objects.

The subjects were first shown the target and then asked to select the item out of the remaining four that was most likely to belong to the same kind of machine/animal as the target. They were then asked to choose the item that was next most likely to belong to the same kind of machine/animal. At the end of the study, the subjects were taken back through each trial and asked why they had made their first and second choices; see Kelemen (1992) for discussion of the self-report data.

Results and Discussion

Preliminary analyses, as well as the self-report data, indicated a strong order effect for the second two trials; as a result, only Trial 1 and Trial 2 were used for each subject.[1] Three 2×2 (kind: natural kind vs. artifact × first target viewed: blue vs. green) analyses of variance

(ANOVAs) were performed. In each case, the dependent variable was the number of times a particular circle (ID[C], ID[S], or OS) was selected at first choice. As predicted, subjects chose the color dimension more often with animals than with machines [$F(1,32) = 4.80, p < .05$]. Subjects also chose the size dimension more often with machines than with animals, although this was only marginally significant [$F(1,32) = 2.21, p = .08$] (see Table 1). No predictions were made about differences in the relative proportion of OS responses, and no significant differences occurred; there was also no effect of first target viewed and no interaction. These findings suggest that background knowledge had affected subjects' classification behavior: color was deemed more important for animals; size was (to some extent) more important for machines.

All second choices were analyzed contingent on first choices. When shown the stimuli described as animals, subjects who gave the ID(C) as a first choice chose the OS as their second choice 82% of the time, instead of the ID(S), which was chosen only 18% of the time. Similarly, subjects who were shown the stimuli described as machines and who chose ID(S) as their first choice made a second choice of the OS 85% of the time, instead of choosing the ID(C), which was selected 15% of the time (both significant by binomial probability; $p < .01$). This suggests that identity does not have intrinsic dominance in a classification task; instead, there is a demand to create a coherent "well-formed" category.

Subjects who chose the OS as their first choice (and disregarded both identity choices) nevertheless conformed to the predicted direction in their second choice, again suggesting the influence of background knowledge. (Animals, 100% color identity; machines, 75% size identity vs. 17% color identity; both significant by binomial probability, $p < .05$.)

A surprising finding has to do with the few subjects who chose the unexpected dimension in their first choice (color for machines and size for animals, rather than color for animals and size for machines). Subjects who made these first choices did *not* show a significant bias to choose the OS as their second choice. This suggests that subjects who picked the expected item as their first choice and then chose OS as second choice did so because they were attempting to form a coherent category—not just because OS was most perceptually similar to their first choice. When subjects were (for whatever reason) not responding in a manner consistent with

Table 1
First Choices for Study 1

Strategy	Natural Kinds	Artifacts
ID(C)	53%	22%
ID(S)	22%	41%
OS	25%	38%

Note—ID(C), item identical to target in color but very different in size; ID(S), item identical to target in size but very different in color; OS, item similar to target in color and size and "overall similar" to the target, relative to other items.

the background information for their first choice—and thus perhaps not forming a conceptually coherent category—their tendency to give the OS response was at chance.

STUDY 2

In the account defended here, categories emerge from the interaction between background knowledge and perceptual information; they are not abstractions from perceptual experience. Given this, one would expect to find the same effect of background knowledge in cases where the stimuli are not presented perceptually at all. To explore this, we attempted to replicate the findings of Study 1 by presenting the stimuli solely through linguistic descriptions.

Method

Subjects. The subjects were 32 University of Arizona undergraduate and graduate students, none of whom participated in the first study.

Materials and Procedure. Each subject was given a booklet that included the instructions that were presented verbally in Study 1, and four trials constructed in the same manner. Instead of seeing the actual objects, however, these subjects were given written descriptions. On one page, the subjects would be given a description of the target object, and on the next page, the target object description would be repeated and they would also be shown the four choices; an example is shown in Figure 3.[2] Each target was labeled with a different name. Order of presentation for the four targets and comparison choices, order of size and color information in the tables, and assignment of stimulus names were all counterbalanced.

The subjects were provided with the questionnaires in a classroom setting; they were told to take as much time as they needed and were asked not to flip forward. On the final page of the questionnaire, the subjects were asked to return through the trials and, without changing their original choices, to explain the motivation for these choices.

Results and Discussion

The dependent variables in these analyses were once again the number of times that a particular circle (ID(C), ID(S), OS) was selected as first choice. A one-way ANOVA (kind: natural kinds vs. artifacts) was performed for each of the three first-choice selections. As in Study 1, subjects chose the color dimension more often with animals than with machines [$F(1,32) = 4.10$,

	SHAPE	COLOR	SIZE
Sample Dish :	Circular	It has a dark green color	It is 35.78 centimeters in size

	SHAPE	COLOR	SIZE
Dish 1:	Same Shape	Identical in color to the sample	Very different size from the sample
Dish 2:	Same Shape	Very different color from the sample	Identical in size to the sample
Dish 3:	Same Shape	Similar color to the sample	Similar size to the sample
Dish 4:	Same Shape	Different color from the sample	Different size from the sample

Figure 3. Sample target object description with four choices for Study 2.

Table 2
First Choices for Study 2

Strategy	Natural Kinds	Artifacts
ID(C)	20%	0%
ID(S)	1.5%	22%
OS	77%	75%
DIST	1.5%	3%

Note—ID(C), item identical to target in color, but very different in size; ID(S), item identical in size to target but very different in color; OS, item similar to target in both color and size and "overall similar" to target, relative to other items; DIST, distractor very different from target in both color and size.

$p < .05$] and chose the size dimension more often with machines than with animals, this time significantly so [$F(1,32) = 4.60, p < .05$].

Unlike in Study 1, however, there was a very high proportion of OS choices (77% for animals; 75% for machines) (see Table 2). The self-report data suggest that this may be an artifact of the wording of the instructions, where the phrase "similar . . . similar" may be viewed as preferable to "identical . . . very different" or "very different . . . identical" and might serve as a cue to some subjects that this choice was in some sense favored by the experimenter. It should be noted, however, that this finding could also be taken as support for the empiricist position; perhaps identity relations expressed through language are less salient for the purposes of categorization than those that are actually presented to the perceptual systems.[3]

As in Study 1, when subjects made their first choices in the expected directions (color for animals, size for machines) they overwhelmingly chose OS for their second choices, opting for a cohesive category as opposed to identity in a different dimension: 100% of the time with animals, 93% for machines (both significant by binomial probability; $p < .01$). In addition, subjects who chose OS for their first choice went in the expected directions for their second choice, attending to color rather than size with animals (animals: 59% color identity vs. 27% size identity) and size rather than color with machines (machines: 60% size identity vs. 17% color identity; both significant by binomial probability $p < .01$).

In sum, the main difference between this study and Study 1 was the high proportion of OS responses as first choices, a finding that can be explained as due either to properties of the instructions or to a more fundamental difference between visually presented arrays versus descriptions of arrays. In all other regards, the results were identical: Color was preferred over size for animals; size was preferred over color for machines; second choices tended to be OS; and the subjects who chose OS for their first choices chose dimensional identity in the expected direction for their second choices.

GENERAL DISCUSSION

The preceding results suggest that the portrayal of classification as a bottom-up perceptual process is an oversimplification, even for stimuli as simple as colored circles. Regardless of whether the stimuli are

actually presented or linguistically described, background knowledge determines which perceptual dimension is relevant. Once a choice has been made, the demand that a category be cohesive dictates future choices.

The motivation for this study was to explore the nature of the developmental differences found in experiments such as those by Smith (1989)—but one might argue that our findings cannot directly address this issue. After all, in those experiments, subjects were shown simple stimuli and asked to sort them, without any explicit context provided. Even if describing stimuli in different ways can affect categorization (as found above), this could not explain the results of Smith, since in those tasks the only information available to the subjects was perceptual, there was no context at all, and thus background knowledge was irrelevant.

This assumption about the nature of such tasks is almost certainly mistaken, however. Adults *do* have knowledge that they can exploit when participating in sorting or categorization experiments, knowledge that children are likely to lack. In particular, adults are aware that the stimuli are artifacts created by a psychologist and that their task is to sort them in the "correct" manner (i.e., in the fashion intended by the psychologist). One implication of this is that any accidental variance in the size or color of the stimuli should be minimal. For each triad used by Smith (1989), for example, two out of the three stimuli (A and B) shared an identity relationship. Adults were likely to have inferred that the existence of such pairs was the creation of the experimenter—not an accident—and it is not surprising that they so consistently obliged by classifying according to this relationship.

Similarly, children's lack of focus on identity in a task such as Smith's (1989) might be the result of their ignorance as to how stimuli are created, a lack of understanding of the intent of the experimenter, and so on. In particular, it might not have occurred to 2- and 3-year-olds that the sets of stimuli were carefully created, with little or no accidental variation, and that the triads were designed so that they would consistently share some common basis for grouping. Children might have chosen the OS item, then, not because of some inability to focus on specific dimensions, but because of ignorance about the nature of psychology experiments.

In support of this, when adults are given speeded classification tasks, in which there is no time to access background knowledge or apply conscious strategies, they also tend to ignore identity and focus on overall similarity (Kemler, 1983). Furthermore, when they are given incidental learning tasks, in which they are not explicitly told to sort the items into groups, adults again adopt a holistic, or overall similarity, mode of classification (Kemler Nelson, 1984). These findings are consistent with the view that adult performance on simple categorization tasks is mediated by background knowledge. When this knowledge cannot be used—either because of time pressure or because the subjects are not explicitly told about the nature of the experiment—they behave like children.

These findings are also consistent with the empiricist explanation of developmental shifts: Classifying by dimension requires attentional resources—children have very limited resources; adults have more, but need time and intent in order to use them. Other considerations, however, militate in favor of the background knowledge theory. For one thing, adults do not *always* sort on the basis of a single dimension, even in intentional nonspeeded tasks. Medin, Wattenmaker, and Hampson (1987) found that for sets of stimuli that have correlated attributes—particularly those that have salient causal linkages between such attributes—adults often sort them into family-resemblance categories; they do not focus on single dimensions. Medin et al. argue that such categories, as opposed to those with uncorrelated attributes, are consistent with adults' expectations that the perceptual properties of an entity are due to "some core factor or cause" (p. 273). Such categories are likely to stand out as the "right" answers in a categorization task (i.e., they stand out as nonrandom intentionally created categories), apparently even more so than those organized on the basis of a single dimension. In contrast, it is unlikely that this finding can be attributed to a failure of attentional focus.

Furthermore, young children have no problem at all focusing on a single dimension when learning the meaning of a new word. When presented with a count noun describing a novel object (e.g., "This is a dax"), 2-year-olds will generalize the word to describe other objects of the same shape, and ignore the dimensions of color, substance, and size (Landau, Smith, & Jones, 1988). This falsifies any claim that children have an across-the-board inability to focus on specific perceptual dimensions in the course of categorization. Most likely children behave this way because they know what word learning is about—more specifically, they know that count nouns used to describe objects refer to kinds of objects and they believe that shape is important when one is classifying some kinds of objects (for discussion, see Bloom, 1994; Landau et al., 1988; Soja, Carey, & Spelke, 1991).

To sum up, in the present experiments adults categorized simple perceptual stimuli differently as a function of how the stimuli were described and they were influenced on subsequent choices by the demand to create "cohesive" categories. This effect occurred both when the stimuli were presented visually and when they were linguistically described. This suggests that children's failure to focus on specific perceptual dimensions in the same kind of task might be due to limitations in background knowledge, not to some lack of perceptual or attentional capacities. This hypothesis is supported by evidence that under some circumstances (speeded classification, incidental learning, exposure to certain types of family resemblance categories) adults act like children and under some circumstances (learning the meanings of words) children act like adults.

REFERENCES

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

BLOOM, P. (in press). Theories of word learning: Rationalist alternatives to associationism. In T. K. Bhatia & W. C. Ritchie (Eds.), *Handbook of language acquisition*. New York: Academic Press.

CAREY, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

EVANS, P., & SMITH, L. B. (1988). The development of identity as a privileged relation in classification: When very similar is just not similar enough. *Cognitive Development*, 3, 265-284.

GELMAN, S. A., & MARKMAN, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, 59, 876-887.

KEIL, F. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

KELEMEN, D. (1992). *The effects of domain-specific knowledge on categorization judgments*. Unpublished master's thesis, University of Arizona.

KEMLER, D. G. (1983). Holistic and analytic modes in perceptual and cognitive development. In T. Tighe & B. E. Shepp (Eds.), *Perception, cognition and development: Interactional analyses* (pp. 77-102). Hillsdale, NJ: Erlbaum.

KEMLER NELSON, D. G. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning & Verbal Behavior*, 23, 734-759.

LANDAU, B., SMITH, L. B., & JONES, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 19-32.

MEDIN, D. L., WATTENMAKER, W. D., & HAMPSON, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, 19, 242-279.

McCLELLAND, J. L., & RUMELHART, D. E. (1986). A distributed model of human learning and memory. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models* (pp. 170-215). Cambridge, MA: MIT Press.

MURPHY, G. L. (1993). Theories and concept formation. In I. van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 173-200). New York: Academic Press.

NISBETT, R. E., KRANTZ, D. H., JEPSON, C., & KUNDA, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, 90, 339-363.

PAZZANI, M. J. (1991). Influence of prior knowledge on concept ac-

quisition: Experimental and computational results. *Journal of Experimental Psychology*, **17**, 416-432.

QUINE, W. V. O. (1960). *Word and object.* Cambridge, MA: MIT Press.

RIPS, L. J. (1989). Similarity, typicality and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). New York: Cambridge University Press.

SMITH, L. B. (1989). A model of perceptual classification in children and adults. *Psychological Review*, **96**, 125-144.

SOJA, N. N., CAREY, S., & SPELKE, E. S. (1991). Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition*, **38**, 179-211.

NOTES

1. Self-reports indicated that the order effect occurred because the animal/machine names used in Trials 1 and 2 were also used in Trials 3 and 4. The intent was that such labels would be construed as superordinates, but instead subjects thought that the labels named the specific stimuli. Thus, in Trials 3 and 4, the subjects made choices that were based directly on their previous choices with the target in Trial 1 or 2 that had the same name.

2. In a pilot study, subjects were given precise size and color descriptions as their choices, instead of phrases such as "similar size to the sample." Many subjects found the amount of information overwhelming in such circumstances and some refused to complete the task. As a result, we shifted to the simpler format shown in Figure 3.

3. We are grateful to Mark McDaniel for bringing this to our attention.