Research Article

LEARNING TO RECOGNIZE OBJECTS

Linda B. Smith

Indiana University

Abstract—*A theory of object recognition requires a theory of shape. Despite considerable empirical and theoretical research, however, a definition of object shape has proved elusive. Two experiments provide new insights by showing that children's object recognition changes dramatically during the period between 17 and 25 months. During this time, children develop the ability to recognize stylized three-dimensional caricatures of known and novel objects. This ability is linked to the number of object names in children's vocabularies, suggesting that category learning may be a driving force behind the developmental changes.*

The central problem for a theory of object recognition is a theory of shape. Although objects seem to be recognized by their shape (Biederman, 1987; Edelman & Duvdevani-Bar, 1997; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976), real instances of real categories are, in fact, rarely the exact same shape. For example, rocking chairs, stuffed chairs, and desk chairs are the "same shape" only under some highly abstract description. Thus, a definition of shape is the central business—and a primary area of contention—in the fields of human and machine object recognition (e.g., Edelman & Duvdevani-Bar, 1997; Hummel, 2000). The experiments reported here provide new insights into this issue by showing that an abstract description of shape—one under which rocking chairs and desk chairs are the same—develops as a product of very young children's category learning.

A COMMON IDEA AMONG COMPETING ACCOUNTS

According to Biederman's (1987; Hummel & Biederman, 1992) recognition-by-components (RBC) account, objects are represented in terms of a set of simple geometric components called geons such that common objects are readily recognized given only two to three geons in the proper spatial arrangement. By this account, the reason a variety of different chairs are seen as chair shaped is that they all conform to the same abstract and componential representation. The RBC account makes no explicit claims about the developmental origins of these representations.

Edelman and his colleagues (Edelman, 1995; Edelman & Duvdevani-Bar, 1997) offer a competing account that begins with the premise that perceivers store view-dependent images of objects (Tarr, 1995). Category learning creates prototypes for (a perhaps small number of) initial categories by interpolating multiple views of multiple instances. Duvdevani-Bar and Edelman (1999) demonstrated that once a number of such prototypes have formed, they serve as landmarks (basis functions) in the shape space such that new objects are represented in terms of their (weighted) distance to those landmarks. In this way, the prototypes for well-learned categories determine the dimensions of the shape space and the perception of even novel things. Although this account differs in important and fundamental ways from

RBC, it is like RBC in that it posits shape representations that are an abstraction over the detailed and specific shapes of real things. Again, a variety of chairs can be seen as chair shaped because they are all highly similar under a particular description of the shape similarity space. Edelman's account also offers a clear-cut developmental hypothesis: Shape similarity should change as a product of early category learning.

THE EXPERIMENTAL QUESTION

The experiments that follow examined this developmental hypothesis. They did not, however, examine the precise nature of the underlying representations, the main point of contention in the adult and machine literature. The experiments examined very young children's recognition of three-dimensional caricatures of the shapes of common things, such as those illustrated at the top of Figure 1. Recognition of impoverished and stylized forms implies a sparse and thus category-encompassing description of shape. When in the course of learning early object categories does such an abstract description of shape emerge? The children participating in the experiments were between the ages of 17 and 25 months. This is a period during which children progress from knowing the names of few objects to knowing the names of several hundred objects (e.g., Bloom, 2000) and a period during which attention to shape in naming tasks increases dramatically (e.g., Smith, 1999).

EXPERIMENT 1

Method

Participants

Twenty-six children (12 male, 14 female) participated in Experiment 1. Their knowledge of object categories was measured by the number of object names in their productive vocabularies. Productive vocabulary was assessed using the MacArthur Communicative Development Inventory (Fenson et al., 1994). Object names were defined as count nouns that referred to concrete entities. The children were divided into two groups according to the number of object names they knew: fewer than 100 object names versus 100 or more object names. Table 1 provides information on the ages and productive vocabularies of the children in each group. The extent of variation in vocabulary among the participants is within the range expected of typically developing children during this age period (see Fenson et al., 1994). However, past research shows vocabulary to be a better predictor of category knowledge than age at these early stages of development (e.g., Waxman, 1998; Xu, 1999).

Stimuli

Lifelike toys in 16 common categories were purchased: hammer, boat, apple, banana, carrot, lollipop, chair, camera, basket, butterfly,

Address correspondence to Linda B. Smith, Department of Psychology, 1101 East 10th St., Indiana University, Bloomington, IN 47405; e-mail: smith4@indiana.edu.

Linda B. Smith

Fig. 1. Photographs of three of the lifelike objects (bottom row) and their caricatures (top row) used in Experiment 1.

cat, pizza, ice cream cone, cake, telephone, and toothbrush. These life-like replicas ranged in size from 10 cm³ to 18 cm³. For each replica, a three-dimensional caricature was constructed from two to four simple geometric components, as illustrated in Figure 1. These components were carved from Styrofoam and painted gray. Although the carica-tures were constructed from simple geometric components in a man-ner consistent with RBC and are described in terms of those components, this in no way implies that recognizing these objects requires parsing them into parts (a central claim of RBC). In brief, the geometric de-scription conforms to how the objects were made but not necessarily to how they are perceived.

Procedure and design

Each child was tested with both lifelike and caricature objects. For each child, eight of the object categories were randomly assigned to the lifelike condition and eight were assigned to the caricature condi-tion such that no child saw both the lifelike and the caricature version of the same object.

Object recognition was measured in a nonlinguistic play task and in a name-comprehension task. The play measure uses the fact that young children spontaneously act on objects in category-specific ways (e.g., McCune-Nicolich, 1981). On each trial of the play task, three unnamed objects were placed in front of the child for 1 min, and the child was encouraged to play ("Look at these. What can you do with these things?"). The child was credited with recognizing an object if the child acted on the object in a category-specific way (e.g., in the case of the phone, pretended to talk on it). On each trial of the name-comprehension task, three objects were placed on a tray. While hold-ing the tray away from the child, the experimenter asked the child to indicate one named object (e.g., "Where is the phone? Show me the phone."). After the question was asked, the tray was moved forward so that the child could respond.

Table 1. *Ages and numbers of object names in the productive vocabulary of the children participating in Experiment 1*

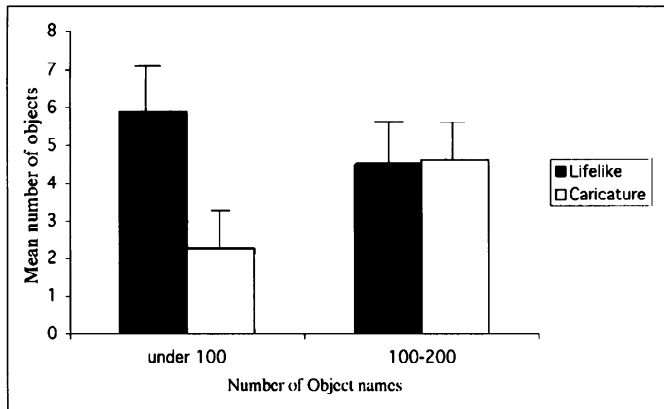| Developmental indicator	Children knowing less than 100 object names (n = 13)		Children knowing 100 or more object names (n = 13)	
	M	Range	*M*	Range
Age (months)	19.5	18.0–21.1	22.0	19.7–24.6
Number of object names	51.6	11–87	173.4	111–237

Object Recognition

Fig. 2. Mean number of lifelike and caricature test objects played with in category-specific ways in Experiment 1. Results are shown separately for children with fewer than 100 and 100 or more object names in their productive vocabulary.

There were 16 trials, each consisting of a 1-min play period (with three objects) and one name-comprehension question (using the same three objects, with one the labeled target and the other two the distractors). All targets and distractors were selected from the set of 16 objects. Half the children received the 8 lifelike trials first and half received the 8 caricature trials first. Position of the correct choice on each trial was randomly determined for each child.

Videotapes of the experimental sessions were coded by a scorer blind to the hypotheses and purposes of the experiment. Six children's videotapes were scored by a second scorer. Agreement for all coded behaviors—category-specific and nonspecific actions during the play period and choice in the name-comprehension task—exceeded .86.

Results and Discussion

As shown in Figure 2, the children with smaller vocabularies played with the lifelike objects in ways that clearly showed they recognized them for what they were, but rarely played with the caricatures in category-specific ways. Apparently, these children with less advanced object-name vocabularies did not recognize the stylized forms as members of the target categories. In marked contrast, the children with larger vocabularies played with the lifelike objects and the caricatures in the same way; they pretended to answer the realistic phone and the shape caricature of a phone; they pretended to eat the realistic plastic pizza and the caricature of the pizza; they made the realistic stuffed cat and the caricature cat meow; they brushed their teeth with the realistic and caricature toothbrush. For these children, the minimalist shape information of the caricatures was sufficient to elicit actions characteristically elicited by the kind.

These conclusions are supported by a 2 (vocabulary level) $\times$ 2 (stimulus object: lifelike vs. caricature) analysis of variance on the number of objects with which children played in category-specific ways. The analysis yielded reliable main effects of vocabulary, $F(1, 24) = 7.69$, $p < .02$, and stimulus object, $F(1, 24) = 8.30$, $p < .01$, and also an interaction between vocabulary and stimulus object, $F(1, 24) = 12.40$, $p < .01$. These conclusions are also supported by the fact that both groups of children interacted with all the objects. Figure 3 shows the number of objects that the children played with in category-specific and nonspecific (stacking, rolling, throwing, showing) ways
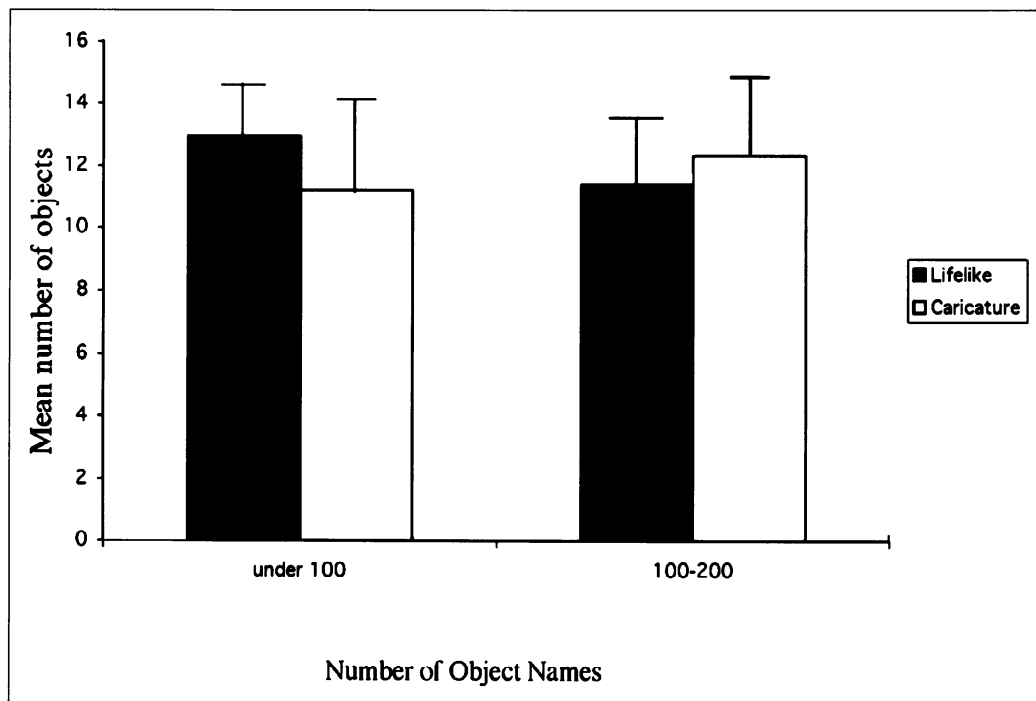

Fig. 3. Mean number of lifelike and caricature test objects acted on (both nonspecific and category-specific actions) in Experiment 1. Results are shown separately for children with fewer than 100 and 100 or more object names in their productive vocabulary.

Linda B. Smith

Fig. 4. Mean number of correct choices of target objects in the name-comprehension task of Experiment 1. Results are shown separately for lifelike and caricature test objects and children with fewer than 100 and 100 or more object names in their productive vocabulary.

combined. The children with smaller vocabularies played with the caricatures as much as did the children with larger vocabularies; they just did not play with the caricatures in category-specific ways.

Figure 4 shows the children's performance in the name-comprehension task. By this measure, the children who knew few object names again recognized the lifelike versions but not the caricatures, and the children with larger vocabularies recognized the lifelike and caricature versions equally well. Again, the more advanced children seemed to need only a stylized form to recognize an object. These conclusions are supported by an analysis of variance that yielded reliable main effects of vocabulary, $F(1, 24) = 66.64, p < .001$, and stimulus object, $F(1, 24) = 23.32, p < .001$, and a reliable interaction between vocabulary and stimulus object, $F(1, 24) = 5.08, p < .05$.

The results provide two new insights into the development of object recognition. First, young children with relatively limited vocabularies do not recognize shape caricatures at all. This tells us that the abstract representations thought to underlie adult object recognition are developmental products. Second, young children who are only slightly more advanced in their category knowledge recognize shape caricatures nearly perfectly, that is, as well as they recognize highly detailed instances. This suggests that the processes of object recognition change rapidly in the same time period that children's knowledge about common object categories expands.

EXPERIMENT 2

Children could learn to represent object shape category by category. Alternatively, they may learn something more general about how to represent shape. Experiment 2 provides preliminary evidence on this issue. Children were taught names for unfamiliar things, and then their recognition of caricatured versions of those things was tested.

Method

Participants

The participants were 60 children (30 males, 30 females) who were divided into two developmental groups (30 children in each group) by their productive vocabulary as in Experiment 1. At each developmental level, children were randomly assigned to one of two conditions: caricatures or shape controls. Table 2 provides information on the ages and productive vocabularies of the children in the experiment.

Table 2. *Ages and numbers of object names in the productive vocabulary of the children participating in Experiment 2*

Developmental indicator	Children knowing less than 100 object names		Children knowing 100 or more object names	
	M	Range	M	Range
	Caricatures			
Age (months)	20.2	18.3–23.4	21.3	17.8–25.0
Number of object names	38.5	10–85	162.2	100–198
	Shape controls			
Age (months)	19.3	17.9–25.5	23.5	17.3–25.5
Number of object names	40.5	10–72	159.3	100–206

Object Recognition

Stimuli

Eight richly detailed (real or toy) objects were selected: artichoke, reamer, masher (potato), manatee, jellyfish, doily, igloo, and waterer (metal watering device). Prior pilot testing of eight 24-month-olds (who did not participate in this experiment) indicated that they did not know the names for these things (mean proportion correct out of a maximum of 8 was .24). For each child in the main experiment, four of these objects were designated as targets and four were designated as distractors, creating four target-distractor pairs. These eight detailed objects served as the starting point for constructing two additional sets of eight. For the caricature condition, as in Experiment 1, a three-dimensional stylized version of each object was constructed from two to four simple geometric shapes carved from Styrofoam and painted gray. Photographs of the original artichoke and its caricature are shown in Figure 5. For the shape-control condition, a duplicate set of the original objects was painted gray such that these shape-control objects were the exact same richly detailed shapes as the originals but like the caricatures differed from the originals in color and texture.

Procedure and design

The task was based on one used previously by Woodward and Hoyne (1999). The events involved are listed in Table 3. In the first training phase, the child was taught the name of one lifelike object. For example, the child might be shown a richly detailed toy artichoke and told, "Look at this artichoke! Wow! See this artichoke? Look—artichoke!" The child was then handed the object to examine. This target object was then removed, and the experimenter introduced the child to the lifelike version of the distractor. The training procedure was the same as with the target object except the distractor object was not named. These two phases of training with each target-distractor pair were repeated three times. The order of the target and distractor phases of training was counterbalanced across children.

The test phase for each of the four target objects consisted of four trials: two memory trials and two experimental trials. The test trials were structured as in the name-comprehension task in Experiment 1 except that the child chose between two objects on each trial, the target and the distractor. On the two memory trials, the target and the distractor were the lifelike versions used in training. The memory trials provide a measure of whether the child linked the name to the original object and also remembered that link. On the two experimental trials,


Fig. 5. Photographs of the realistic artichoke and its caricature used in Experiment 2.

the target and the distractor were the caricature or shape-control versions of the lifelike objects, depending on the condition to which the child was assigned. For each named target, the experimental trials were always the second and third test trials.

This procedure of training and then test was repeated for each of the four target-distractor pairs, for a total of eight memory trials and eight experimental trials. Position of the correct choice on the test trials was randomly determined, with each target appearing equally often on the right and on the left.

The children's performance in the shape-control condition provides information about children's potential difficulty in recognizing the shape caricatures. Poor performance on the caricature tests could be due to an inability (or refusal) to generalize a name to an object that differs in any way from the trained target, rather than to an inability to recognize stylized shapes. If children's name extensions in the caricature condition were limited only by their lack of sufficiently abstract shape representations, then children would readily extend targets' names to the shape controls because these objects matched the trained targets perfectly in shape.

A coder blind to the hypotheses coded the videotapes of the sessions. A second coder scored the responses of 4 children. Correspondence for the two coders was 100%.

Results and Discussion

Overall the children performed well on the memory trials, choosing the labeled target on 75% of the trials. A 2 (vocabulary level) × 2

Table 3. *Phases of the procedure in Experiment 2*

| Phase	Experimenter's language	Stimuli present
Training—target	Look at this artichoke. Wow! See this artichoke? Look—artichoke.	A richly detailed and lifelike artichoke
Training—distractor	Look at this. Wow! See this. Look.	A richly detailed and lifelike jellyfish
Memory test	Where is the artichoke?	Choice between the richly detailed artichoke and the richly detailed jellyfish
Experimental test	Where is the artichoke?	Choice between the shape caricature of the artichoke and the shape caricature of the jellyfish

Linda B. Smith
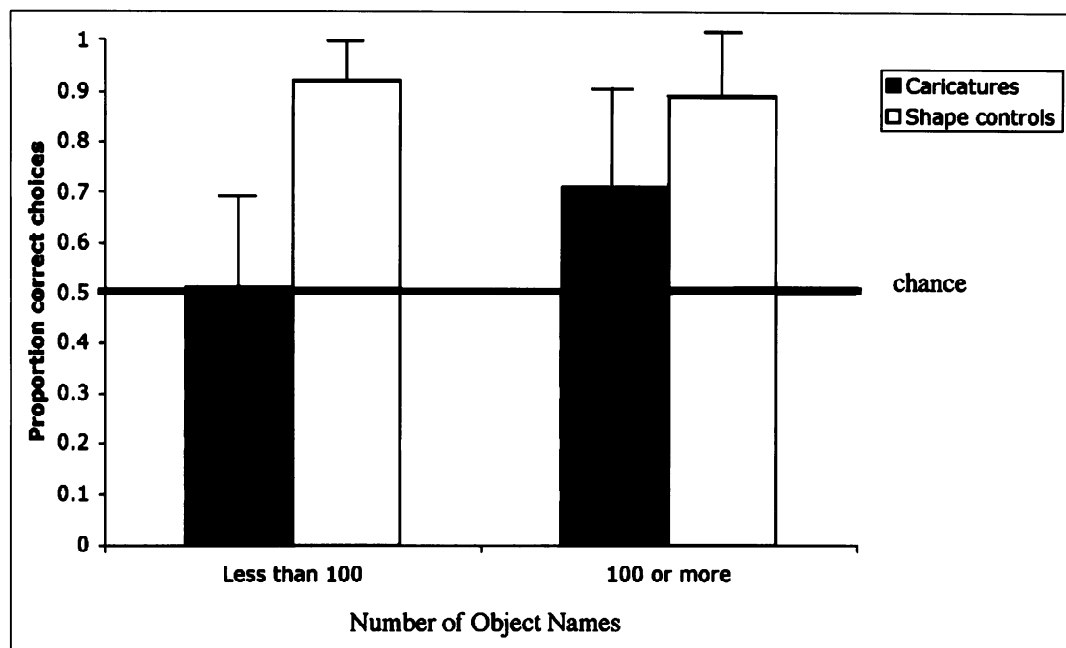
Fig. 6. Proportion correct on the experimental trials (given correct performance on the memory trials) in the caricature and shape-control conditions of Experiment 2. Results are shown separately for children with less than 100 and 100 or more object names in their productive vocabulary.

(condition) analysis of variance of correct choices on the memory trials revealed no main effects or interactions, $p > .36$ in all cases. Thus, the children with smaller and larger vocabularies did not differ in their ability to remember the target objects' names.

Performance on the experimental trials was assessed in terms of a conditional probability: the probability of a correct choice on the experimental trials given that the child had selected the target correctly on both memory trials for that object. The results shown in Figure 6 indicate that whereas all children recognized the shape controls, only the children with more advanced object-name vocabularies recognized the caricatures. This was confirmed by a 2 (vocabulary level) $\times$ 2 (condition) analysis of variance, which yielded a reliable main effect of condition, $F(1, 54) = 58.9, p < .001$, and a reliable interaction between vocabulary and condition, $F(1, 54) = 9.79, p < .01$. Post hoc analyses (Tukey's HSD, $p < .05$) indicated that the children with more advanced vocabularies extended names to the caricatures more than did the children with less advanced vocabularies. These results indicate that in this period of rapid growth in object naming and categorization, children are learning something general about category-relevant definitions of object shape.

GENERAL DISCUSSION

As children's knowledge of common object categories expands, children's perception of shape similarity also changes. Experiment 1 suggests that during this period, children determine the bare-bones essentials of shape that make a thing cup shaped or camera shaped or cat shaped. Experiment 2 suggests that children also learn how to represent the shapes of even novel things. These results are consistent with the role of category learning in Duvdevani-Bar and Edelman's (1999) computational model. The results also fit well with a growing body of findings suggesting that perception itself may be molded by and not be separable from top-down processes and category learning (Goldstone & Barsalou, 1998; Goldstone, Lippa, & Shiffrin, 2001).

Although infants' perception of three-dimensional object shape has not been extensively studied (see Kellman, 2001), there is considerable evidence of very early sensitivities to shape and even abstract shape similarities. For example, infants perceive the similarities between richly detailed three-dimensional objects and their two-dimensional (photographic) representations (DeLoache, Strauss, & Maynard, 1979), they abstract shape similarities across unique exemplars in habituation studies (e.g., Bomba & Siqueland, 1983), and they are sensitive to the invariant aspects of the shape of a single object from multiple perspectives (e.g., Slater & Morison, 1985). These abilities are in place before category learning and presumably contribute to the early learning of object categories. The present results suggest, however, that learning object categories changes (or tunes) the way children perceive shape, creating descriptions of object shape that emphasize shape properties crucial to object recognition and that deemphasize the less relevant aspects of shape. In this way, kitchen chairs, rocking chairs, and chair caricatures come to be the same shape.

In its full form, the developmental story behind mature object recognition is likely to be a long one, dependent on extensive category experience. The results from Experiment 1 hint that this is the case. In that experiment, the children with the most limited vocabularies readily recognized the lifelike and richly detailed instances, and indeed did so as well as the children with more extensive vocabularies. Yet the children with the smaller object-name vocabularies did not recognize the caricatured shapes of these things. Apparently, this ability emerges well after children have considerable proficiency with many individual categories.

Studies of categorization in infancy also suggest a long developmental course. Infants as young as 3 and 4 months categorize, for example, dogs as different from cats (e.g., Eimas & Quinn, 1994; Quinn, Eimas, & Rosenkrantz, 1993). However, a compelling series of experimental results suggests that they do not use overall shape (see Quinn, in press). Experiments with older infants support this idea. Xu and Carey (1996; Carey & Xu, 2001; Xu, in press) found that 10- to 12-month-old infants seem not to remember the *kind* of thing that they have just seen. For example, if infants see a duck disappear behind a screen, they are not surprised to see a truck appear, although they are surprised to see two ducks (or two trucks) appear. It is as if infants remember the number of things but not their shape. However, infants with larger receptive vocabularies are more likely to recognize just-seen objects in this task than are less advanced infants. Further, their ability to distinguish kinds appears closely related to their knowledge of adult lexical categories. For example, although a two-handled baby cup, a one-handled conventional cup, and a baby bottle are all roughly similar in shape but all different from one another in details, infants distinguish the cups from the bottle but not one cup from another, a pattern in line with conventional lexical categories. Altogether, these results suggest a progressive definition of shape, one that may literally be created through learning object categories.

The present findings also may contribute to our understanding of other developmental changes that occur in this same time frame. For example, developing descriptions of object shape may help us explain how children generalize the functions of objects in principled ways (Gelman & Bloom, 2000; Kemler Nelson, Russell, Duke, & Jones, 2000; Smith, in press). Changes in shape perception may also contribute to the emergence of symbolic play, that is, play in which one object is creatively substituted for another. For example, young children have been reported to use a banana as a phone, a shoe box as a bed, or a stick as a bottle (Corrigan, 1982; Mc-Cune-Nicolich, 1981; Shore, O'Connell, & Bates, 1984). Symbolic play substitutions—for example, using a banana as a phone—require children to perceive sparse shape similarities.

In conclusion, these experiments are a first step in the developmental study of object recognition and thus leave unanswered many questions, including questions about the particular experiences crucial to these developmental changes and the role of language learning. The present results also leave unspecified the exact nature of the object representations that children build, and thus do not distinguish between object- and image-based accounts of object recognition. However, the results do provide a hint about shape representations that might be relevant. The shape caricatures used in these experiments differed substantially from real objects in that they included none of the details of real object shapes and few of the parts. Thus, the caricatured cat had no eyes nor tail, but only two triangles set on a sphere set on a rectangular block. The caricatures preserved overall shape by presenting a small number of simple geometric components in their proper spatial arrangement, just as would be predicted by Biederman's (1987) RBC account. Perhaps there will be a developmental solution to the debate between object-based and image-based theories. Representations of the shapes of things—representations perhaps ultimately describable in terms of simple components—may be created out of view-dependent images as a consequence of category learning.

REFERENCES

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–117.

Bloom, P. (2000). *How children learn the meanings of words.* Cambridge, MA: MIT Press.

Bomba, P.C., & Siqueland, E.R. (1983). The nature and structure of infant form categories. *Journal of Experimental Child Psychology, 35,* 294–328.

Carey, S., & Xu, F. (2001). Infants' knowledge of objects: Beyond object files and object tracking. *Cognition, 80,* 179–213.

Corrigan, R. (1982). The control of animate and inanimate components in pretend play and language. *Child Development, 53,* 1343–1353.

DeLoache, J., Strauss, M., & Maynard, J. (1979). Picture perception in infancy. *Infant Behavior and Development, 2,* 77–89.

Duvdevani-Bar, S., & Edelman, S. (1999). Visual recognition and categorization on the basis of similarities to multiple class prototypes. *International Journal of Computer Vision, 33,* 201–228.

Edelman, S. (1995). Representation, similarity, and the chorus of prototypes. *Minds & Machines, 5*(1), 45–68.

Edelman, S., & Duvdevani-Bar, S. (1997). A model of visual recognition and categorization. *Philosophical Transactions of the Royal Society of London, 352,* 1191–1202.

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.

Fenson, L., Dale, P.S., Reznick, J.S., Bates, E., Thal, D.J., & Pethick, S.J. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development, 59*(5, Serial No. 242).

Gelman, S.A., & Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. *Cognition, 76,* 91–103.

Goldstone, R.L., & Barsalou, L.W. (1998). Reuniting perception and conception. *Cognition, 65,* 231–262.

Goldstone, R.L., Lippa, Y., & Shiffrin, R.M. (2001). Altering object representations through category learning. *Cognition, 78,* 27–43.

Hummel, J.E. (2000). Where view-based theories break down: The role of structure in human shape perception. In E. Dietrich & A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 157–185). Mahwah, NJ: Erlbaum.

Hummel, J.E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–517.

Kellman, P. (2001). Separating processes in object perception. *Journal of Experimental Child Psychology, 78,* 84–97.

Kemler Nelson, D.G., Russell, R., Duke, N., & Jones, K. (2000). Two-year-olds will name artifacts by their functions. *Child Development, 71,* 1271–1288.

McCune-Nicolich, L. (1981). Toward symbolic functioning: Structure of early pretend games and potential parallels with language. *Child Development, 52,* 785–797.

Quinn, P.C. (in press). Young infants' categorization of humans versus nonhuman animals: Rules for knowledge access and perceptual process. In L. Gershkoff-Stowe & D. Rakison (Eds.), *Building object categories in developmental time: 32nd Carnegie symposium on cognition.* Mahwah, NJ: Erlbaum.

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.

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

Shore, C., O'Connell, B., & Bates, E. (1984). First sentences in language and symbolic play. *Developmental Psychology, 20,* 872–880.

Slater, A.M., & Morison, V. (1985). Shape perception and slant perception at birth. *Perception, 14,* 337–344.

Smith, L.B. (1999). Children's noun learning: How general learning processes make specialized learning mechanisms. In B. MacWhinney (Ed.), *The emergence of language* (pp. 277–303). Mahwah, NJ: Erlbaum.

Smith, L.B. (in press). Shape: A developmental product. In L. Carlson & E. van der Zee (Eds.), *Functional features in spatial concepts.* Oxford, England: Oxford University Press.

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

Waxman, S.R. (1998). Linking object categorization and naming: Early expectations and the shaping role of language. In D.L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 38, pp. 249–291). San Diego, CA: Academic Press.

Woodward, A.L., & Hoyne, K.L. (1999). Infants' learning about words and sounds in relation to objects. *Child Development, 70,* 65–77.

Xu, F. (1999). Object individuation and object identity in infancy: The role of spatiotemporal information, object property information, and language. *Acta Psychologica, 102,* 113–136.

Xu, F. (in press). The development of object individuation in infancy. In J. Fagen & H. Hayne (Eds.), *Progress in infancy research* (Vol. 3). Mahwah, NJ: Erlbaum.

Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology, 30,* 111–153.