

Seeing Things From a Different Angle: The Pigeon's Recognition of Single Geons Rotated in Depth

Jessie J. Peissig, Michael E. Young,
and Edward A. Wasserman
University of Iowa

Irving Biederman
University of Southern California

In 2 experiments involving computer-rendered versions of single shapes or "geons," the extent to which depth rotation affects the visual discrimination performance of pigeons in both go/no-go and forced-choice tasks was documented. The pigeons were able to recognize geons at most rotations in depth; however, the pigeons' recognition performance was better at the training view than at most other views. Aspects of these results are both consistent with and problematic for object-centered and viewer-centered theories of object recognition.

As an organism moves about its visual environment, it appears to identify and to categorize objects effortlessly and instantaneously. To do so, an organism must be able to recognize objects from varying distances and viewpoints, despite drastic differences in the object's absolute size and apparent shape.

The process of object recognition requires that an object in the visual field be compared with representations of objects stored in memory. One class of theories proposes that these stored representations are independent of the viewpoints at which they are seen in the environment, thus allowing the organism to recognize objects from many different viewpoints without having to store each view. In these "object-centered" theories, an object in the visual field is compared with three-dimensional representations that are stored in memory. For recognition to occur, the object's three-dimensional structure must then be derived from its two-dimensional retinal image and matched with a stored representation (Biederman, 1987; Corballis, 1988; Marr & Nishihara, 1978); if the derived structure matches a representation in memory, then it is recognized. For example, a cup can be thought of as a cylinder with one closed end. A cup with a handle has a smaller, bent cylinder attached to the side of the main cylinder. A retinal image of the cup with the handle that maintains this basic two-component structure will be recognized as a cup. Even novel views of the cup are recognizable, if the derived structure matches closely with the appropriate representation.

One of the most widely known of the object-centered theories is the recognition-by-components (RBC) theory proposed by Biederman (1987). According to RBC, objects are decomposed into arrangements of simple geometric components called "geons." An object can be recognized from different viewing angles and distances when its geons and their relational positions are consistent from view to view. If, however, rotation were to reveal previously unseen geons or to conceal previously seen ones, then RBC predicts that the accuracy of object recognition will be impaired, require additional time, or both.

For example, Biederman and Gerhardstein (1993, Experiment 3) used nonsense objects comprising one central geon and two pairs of geons attached to the front, back, left, and right of the central geon. People's reaction times were significantly longer when a 45° rotation in depth revealed new parts than when the same nominal rotation did not reveal new parts.

Avian Object Recognition

The majority of nonhuman research in the area of visual object recognition has used monkeys as subjects (Gross, Bender, & Rocha-Miranda, 1972; Logothetis, Pauls, & Poggio, 1995; Tanaka, 1996). There is a growing body of research, however, exploring object recognition in pigeons (Cerella, 1977; Cook & Katz, 1999; Reid & Spetch, 1998).

In one recent pigeon project, Wasserman et al. (1996) examined viewpoint effects on visual recognition using line drawings of complex, multiple-geon objects. Pigeons were trained to discriminate among four objects, each at a single view, by differential reinforcement of pecking to four report keys. In Experiment 1a, three groups of pigeons were trained with each of the four objects shown at a 0°, 33°, or 67° view and they were then tested with the two untrained views. In Experiment 1b, the pigeons were tested with the same objects across a broader range of views (spanning ±133°). In both experiments, the pigeons exhibited statistically significant generalization to other views after training at only one view. The pigeons' accuracy, however, decreased for views farther away from the trained view, thereby documenting viewpoint-specific responding.

Jessie J. Peissig, Michael E. Young, and Edward A. Wasserman, Department of Psychology, University of Iowa; Irving Biederman, Department of Psychology, University of Southern California.

This research was supported by a National Institute of Mental Health Grant MH 473131. We thank Oliver Dreer for supplying the PictCompare program and source code. We also thank B. J. Terrones and T. Barclay for their assistance collecting data.

Correspondence concerning this article should be addressed to Jessie J. Peissig or Edward A. Wasserman, Department of Psychology, University of Iowa, Iowa City, Iowa 52242-1407. Electronic mail may be sent to jessie-peissig@uiowa.edu or ed-wasserman@uiowa.edu.

Viewpoint-specific responding would initially seem to be inconsistent with an object-centered theory of recognition such as RBC. RBC does, however, predict a generalization decrement when geons are revealed and concealed as objects rotate in depth, which occurred with the multiple-component stimuli used in Wasserman et al.'s (1996) project.

Stimulus Considerations

A particularly important factor in studying visual object recognition is the type of experimental stimuli used. Biederman and Gerhardstein (1993) proposed that viewpoint-invariance in object recognition will occur only if the stimulus satisfies three criteria: (a) the object is decomposable into geons, creating a specific geon structural description (GSD) of the geons and their spatial relations, (b)

different objects have distinctive GSDs, and (c) the GSD of each object remains the same across tested views. To critically test RBC, the experimental stimuli must meet these limiting criteria. Single geons (see Figure 1) are the most basic components of objects and therefore, by definition, meet the three criteria set by RBC. They form the simplest GSD, are distinctive from one another, and are recognizable at all viewpoints (with the exception of accidental viewpoints; Biederman, 1987).

Biederman and Gerhardstein (1993, Experiment 4) used line drawings of single geons to explore the effects of rotation in depth on recognition performance. Human participants were shown a single geon for 20 s prior to the start of a block of trials; they were instructed that this geon would serve as the target for that block of trials. The participant's task was to press a key if the stimulus was the target and to

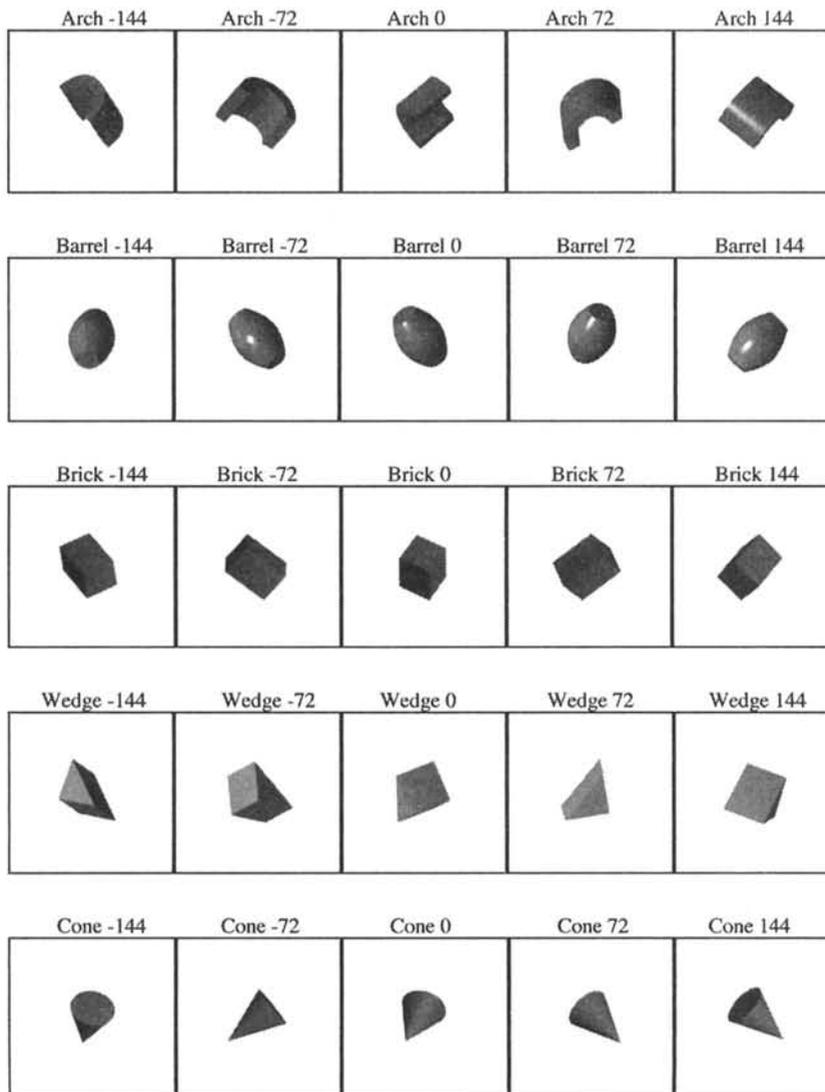


Figure 1. The five geons that served as stimuli in Experiment 1, shown at all views used in this experiment. The geons were arch, barrel, brick, wedge, and cone.

make no response if the stimulus was not the target. Within a block of trials, participants were presented with views of the same geons rotated 45° and 90° in one direction. Biederman and Gerhardstein (1993) found that there were no significant differences in reaction times across the three views. These results suggested that recognition of single geons is unaffected by rotation, supporting RBC's proposition that object recognition is viewpoint invariant.

Hayward and Tarr (1997, Experiment 2), however, reported results that contradict those of Biederman and Gerhardstein (1993). Hayward and Tarr (1997) used a matching task in which the participants were shown two single-geon objects presented sequentially. The participants were instructed to make a response if the two objects were the same as each other and to make no response if the two objects were different from one another. On same-object trials, the participants saw the 0° view of an object paired either with itself or with one of two 45° rotations of the same object. One 45° rotation involved a qualitative change from the 0° view; a qualitative change was defined as a significant change in the "bounding contour" of the object (p. 1517). The remaining 45° rotation involved a quantitative change from the 0° view; a quantitative change was defined as a minimal change in the bounding contour of the object. On different-object trials participants saw the 0° view or one of the 45° rotations paired with one of nine distractor objects. Ten different single-geon objects were used, each of which served as the object on same-object trials and as a distractor on different-object trials.

On same-object trials, participants responded significantly slower when the 0° view of an object was paired with the 45° rotations of that object than when it was paired with itself. Hayward and Tarr (1997) also found participants' performance was significantly slower and less accurate for the 45° rotation involving a qualitative change than for the 45° rotation involving a quantitative change. On the basis of the results of their study, Hayward and Tarr (1997) concluded that single geons have rotational costs, even though they do not violate any of Biederman and Gerhardstein's (1993) limiting criteria, and these costs are greater for a rotation that produces a qualitative change than for a rotation that produces a quantitative change. Hayward and Tarr's (1997) results are contrary to what is expected if geons are viewpoint-invariant building blocks of a generalized structural description, as Biederman and Gerhardstein (1993) have claimed.

It is unclear whether the particular task or methods used in Hayward and Tarr (1997) and Biederman and Gerhardstein (1993) may have influenced the way in which the participants responded. It seems unlikely, however, that the difference in the pattern of results is due only to task differences. Tarr, Williams, Hayward, and Gauthier (1998) performed a study using shaded single-geon images in three different tasks: (a) a sequential matching task identical to that used by Hayward and Tarr (1997), (b) a matching-to-sample task identical to the task used by Biederman and Gerhardstein (1993), and (c) a naming task. Tarr et al. (1998) reported rotation costs in all three tasks, although the effect

was smaller for the matching-to-sample task. Tarr et al.'s (1998) results suggest that a task difference alone cannot explain the conflicting results. Thus, additional studies using single geons may clarify these issues.

In addition to contributing to the current research in object recognition, our study also broadens the range of rotations that have been studied. Previous experiments using single geons have explored rotation in depth across a very limited range of values. In the experiment reported by Biederman and Gerhardstein (1993), the target geon differed in rotation by either 45° or 90° in one direction. In the experiment reported by Hayward and Tarr (1997), geons differed by 45° in one direction of rotation or by 45° in the opposite direction. In the natural environment, the viewpoint at which we encounter objects is rarely so restricted. Increasing the number and range of depth rotations may yield additional insights into the object recognition process.

The Current Project

In our study, we explored how pigeons generalize conditioned responding at one view of a single-geon object to other views of that object. They were trained with only one view of single-geon objects to provide the cleanest examination of viewpoint-invariance. If geons are the basic components of objects, as RBC theory proposes, then training at one view should allow the formation of a consistent GSD. Therefore, despite seeing only a single training view, the pigeons should exhibit viewpoint invariance if they are using an object-centered recognition mechanism. In addition, even though the pigeons were trained with only one view, they were tested with a broad range of rotations in depth. Thus, this study not only explored the nature of object recognition mechanisms in nonhuman animals, but it also examined the issue of recognizing single geons beyond what has been reported in experiments with human participants.

Experiment 1

Studies of object recognition in humans have used sequential matching, matching-to-sample, or naming tasks (e.g., Biederman, 1987; Biederman & Gerhardstein, 1993; Hayward & Tarr, 1997; Tarr, 1995; Tarr, Bühlhoff, Zabinski, & Blanz, 1997; Tarr et al., 1998). The principal dependent measure in these studies is reaction time, although the accuracy of people's reports is of strong interest and importance as well.

With pigeons and other nonhuman animals, reaction times are often an insensitive measure of discriminative behavior; reaction times are particularly problematical when accuracy systematically differs from the training to the testing stimuli (see Wasserman et al., 1996 for further discussion of reaction time measurement in behavioral work with pigeons). The go/no-go paradigm, however, has proven to be a very sensitive method for determining the perceived similarity of a variety of visual stimuli (e.g., Kirkpatrick-Steger & Wasserman, 1996). In the go/no-go paradigm, a reinforced stimulus is displayed for a predetermined amount of time; after the time has elapsed, the subject is required to make a

response for reinforcement to be given (a "go" trial). A nonreinforced stimulus is displayed for the same predetermined amount of time and terminated, regardless of whether or not the subject makes a response (a "no-go" trial). As the pigeon learns to discriminate reinforced from nonreinforced stimuli, an increasing percentage of its responses (pecks) occurs on the reinforced trials. By examining peck rates during discrimination learning, one is able to determine the perceived similarity of each nonreinforced stimulus to the reinforced stimulus.

In the current experiment, we examined the pigeon's discrimination performance when it was reinforced for responding to only one view of a target geon. Pecks to other views of the target geon or to any view of a foil (nontarget) geon were not reinforced. This task provides a simple test of object-centered theories; if object recognition is viewpoint invariant and object differences are more salient than viewpoint differences, then pigeons should respond more to the target geon than to the foil geons, regardless of the view of the target geon.

Method

Subjects. The subjects were 16 feral pigeons (*Columba livia*). By controlled daily feeding, the pigeons were maintained at 85% of their free-feeding weights. Prior to the start of the present experiment, the pigeons had participated in unrelated studies.

Apparatus. The pigeons were trained in four specially constructed plywood chambers. One side of each chamber consisted of a large opening with an aluminum frame attached to the outside of the box. Inside the frame was a clear touch screen (Elmwood Sensors DuraTouch; Model #70056-001, Pawtucket, RI) that was coated with mylar for durability. The pigeon's pecks on the touch screen were processed by a serial controller board (Elographics; Model #E271-2200, Oak Ridge, TN). A brushed aluminum panel was placed directly in front of the screen to allow pigeons access to limited portions of the video monitor. Although there were five openings in the aluminum panel, this experiment used only the central 7 × 7-cm square opening in which the stimuli appeared. In the rear of the chamber, a clear Plexiglas food container was placed level with a wire mesh floor to prevent pigeons from perching on the food cup. Noyes 45-mg pigeon pellets were delivered through a vinyl tube into the food cup using a rotary pellet dispenser (Model #ENV-203M; MED Associates, Lafayette, IN). During experimental sessions, constant illumination was provided by a houselight mounted on the upper rear wall of the chamber. A digital input-output (I/O) interface board (National Instruments; Model #NB-DIO-24, Austin, TX) controlled the pellet dispenser and the houselight.

Control of peripheral stimuli (via the I/O interface) and recording of the pigeon's responses (via the serial controller board) were accomplished by two Apple Macintosh IIfx computers for two of the four chambers. The pigeon's monitor and an identical monitor located in an adjacent room were connected by a video splitter (VOPEX-2M, Network Technologies Inc., Aurora, OH). For the other two chambers, two Apple Macintosh Quadra 650 computers were used in controlling the peripheral stimuli and recording the pigeon's responses. The pigeon's monitor and an identical monitor in an adjacent room were connected by a distribution amplifier (MAC/2 DA 2, Extron Electronics, Sante Fe Springs, CA). The programs were developed in Hypercard Version 2.3.

Stimuli. Geon stimuli were rendered in Ray Dream Studio at 300 dpi resolution. The stimuli ranged from 2 to 4 cm in width and

from 2 to 4 cm in height; they consisted of an arch, a barrel, a brick, a wedge, and a cone that differed from one another by a variety of nonaccidental properties (Biederman, 1987). The geons were rotated in depth by 72° intervals to produce five different views of each geon (see Figure 1 for representations of the pictures that are inferior in quality to the pictures the pigeon actually saw). The "brick wall" negative stimulus was a pattern of line-drawn horizontal 0.29 × 0.12 cm rectangles that filled the entire center display area. It did not contain any shading or depth cues to maximize its discriminability from the geon stimuli.

Procedure. Sixteen pigeons were randomly assigned to one of four conditions: arch+, barrel+, brick+, and wedge+. (To satisfy counterbalancing across the four experimental chambers, the cone was not assigned as a target.) The experiment consisted of two different phases: a preliminary discrimination phase and a final discrimination phase. During preliminary discrimination training, the pigeons were trained with only two stimuli: the reinforced view of the target geon that served as the positive stimulus (S+) and a black-and-white "brick wall" that served as the negative stimulus (S-). We taught the pigeons to respond differentially to a single view of the target geon during preliminary training in order to record any spontaneous generalization that might occur in the final discrimination phase to the other 24 geon stimuli.

At the beginning of a trial, the center display area was illuminated with a white field. A single peck anywhere within that display area turned on the stimulus in the center display area. On a nonreinforced (no-go) trial, the stimulus remained on for 15 s after which the display area darkened and the intertrial interval began; all pecks during the 15 s were recorded. On a reinforced (go) trial, the stimulus remained on until the first peck after 15 s, which delivered reinforcement; the number of pecks during the first 15 s was recorded. The intertrial interval ranged from 6 to 10 s (mean of 8 s). Pigeons were reinforced with 1 to 3 pellets of food, determined for each bird by its weight.

The preliminary discrimination phase consisted of five randomized blocks of 35 trials: 10 go trials with the reinforced view of the target geon (S+) and 25 no-go trials with the brick wall (S-). There were five blocks in each daily session for a total of 175 trials per day. Before moving from preliminary to final discrimination training, all pigeons were required to meet a criterion of at least four consecutive sessions in which the rate of pecking to the nonreinforced brick wall fell to less than 33% of the rate of pecking to the reinforced view of the target stimulus.

The final discrimination phase consisted of five randomized blocks of 35 trials. Within each block, there was a nonreinforced presentation of each of the 5 views of each geon (25 total nonreinforced trials) and 10 reinforced trials with the S+. This arrangement resulted in 24 S- trials and 11 S+ trials (10 reinforced and 1 nonreinforced) in a block. This method of reinforcing 29% of the trials was used to assure that the pigeons maintained responding. Reinforcing 91% of the S+ trials also allowed the pigeons to effectively discriminate the S+ from the S-s, which were never reinforced. The data used for statistical analyses came only from the nonreinforced trials, so that each of the 25 pictorial stimuli contributed equally to the data analysis.

Data collection began on Day 1 of the final discrimination phase to determine if any spontaneous generalization would occur to other views of the target geon or to any views of the foil geons. Data collection ended after the rate of pecking (in a single session) to any 23 of the 24 S-s reached 33% of that to the S+ (this was the best level of discrimination performance that was achieved by all 16 pigeons). To simplify data analysis and interpretation, the five views of the four foil geons for each pigeon were combined to create a single measure of peck rate to the foils.

Results

The pigeons spent a mean of 6.6 days in preliminary discrimination training. The mean number of scored days in final discrimination training was 15.8. For final discrimination training, pigeons in the arch+ condition took the fewest days ($M = 6.5$) to reach criterion and pigeons in the wedge+ condition took the most days ($M = 24.0$) to reach criterion; pigeons in the remaining conditions took an intermediate number of days (barrel+ condition: $M = 18.3$, and brick+ condition: $M = 14.3$) to reach criterion.

We analyzed the first day of final discrimination training to determine if preliminary discrimination training produced differential responding to any views of the target geon. To do so, we calculated the percentage of pecks to the S+, calculated by dividing the rate of pecking to each stimulus by the rate of pecking to the S+ and multiplying by 100. This procedure always yields 100% for the 0° rotation of the target geon (S+). Percentages were used instead of the rate of responding to compare the behavior of different pigeons, which exhibited very different rates of pecking. The percentage of pecks for nonreinforced views of the target geon was 97.5%, compared to 97.0% for the foil geons. Looking at each condition individually, in the arch+ condition, the percentage was 98.6% for the nonreinforced target views and 99.1% for the foils. In the barrel+ condition, the percentage was 103.7% for the nonreinforced target views and 104.2% for the foils. In the brick+ condition, the percentage was 89.0% for the nonreinforced target views and 92.5% for the foils. In the wedge+ condition, the percentage was 98.7% for the nonreinforced target views and 92.4% for the foils. These data indicate that the pigeons did not exhibit differential generalization to different views of the target geon after preliminary discrimination training.

To confirm these initial observations, the data for Day 1 of final discrimination training were analyzed using a full factorial repeated-measures analysis of variance (ANOVA) of condition (arch+, barrel+, brick+, and wedge+) and degrees of rotation from the target (-144°, -72°, 0°, 72°, and 144°), nested within geon class (target vs. foils). The dependent measure was the rate of pecking to each stimulus. The rate of pecking for each stimulus per day was used as the dependent measure instead of percentages, because the repeated-measures ANOVA partitions any individual differences in pigeon peck rates into separate variance components. There were no significant main effects and only one significant interaction. The condition by geon class interaction was significant, $F(3, 12) = 4.42, p < .05$, indicating that the pigeons' responses to the target and foil geons varied between conditions. Post hoc analyses (Newman-Keuls, $\alpha = .05$) indicated that pigeons in the wedge+ condition responded significantly more to the target geon than to the foil geons on the first day. Pigeons in the other three conditions did not exhibit any significant differences in responding. From these data, we can conclude that if preliminary discrimination training generated differential responding, then it was only for pigeons in the wedge+ condition.

Table 1 depicts the mean results of all scored days of final

Table 1
Percentage of Pecks to the S+ in Experiment 1
for Each Bird at Each View by Condition

Condition and bird	Target geon					Foil geons
	-144°	-72°	0°	72°	144°	All views
Arch+						
10B	50	45	100	61	48	48
35R	60	63	100	74	67	69
36R	33	30	100	48	48	37
51R	39	32	100	64	45	37
Barrel+						
25Y	77	97	100	75	52	53
31R	42	62	100	60	39	38
78B	52	99	100	70	35	31
79W	43	83	100	61	51	41
Brick+						
15B	78	70	100	73	86	52
38R	46	43	100	48	65	39
3Y	82	68	100	78	96	64
46R	89	49	100	63	72	31
Wedge+						
27R	38	48	100	95	77	36
33R	18	33	100	100	86	28
49R	27	40	100	99	84	38
88Y	24	38	100	97	83	21
<i>M</i>	50	56	100	73	64	41

discrimination training for each pigeon. The data are presented as the percentage of pecks to the S+, calculated by dividing the rate of pecking to each stimulus by the rate of pecking to the S+ and multiplying by 100. The percentages for views of the target geon are shown individually and the percentages for all the views of the foil geons are averaged together to yield a single mean. Table 1 shows that the pigeons better discriminated nonreinforced views of the target geon as they rotated farther away from the reinforced (S+) view. Table 1 also shows the percentage of pecks across all five views of the four foil geons; across all 16 pigeons, that score was 41%, lower than any mean score registered to the never-reinforced views (-144°, -72°, 72°, 144°) of the target geon.

Figure 2 is a polar plot that more vividly illustrates the overall percentage of responding to the target geon, as compared to responding to the foils, across all four stimulus conditions. The gray area represents the percentage of pecks, averaged across pigeons, to each rotation of the target geon. The point where the gray area intersects the spokes indicates the precise percentage. For example, in Figure 2 at the 72° view, the gray intersects approximately three fourths of the way up that spoke of the graph, indicating that the percentage of pecks is approximately 75% (Table 1 shows that the actual score is 73%). The dashed circle represents the mean percentage of pecks (41%) to all five views of the four foil geons combined. The individual views of the foils were not related in any systematic way to the views of the target geon; therefore, the data for the foils were collapsed across views (see Appendix A for a full table of pigeons' responding to the foils).

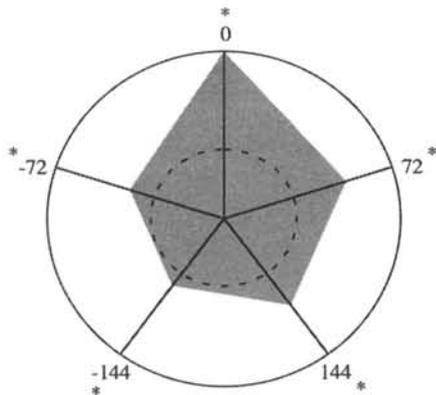


Figure 2. The results of Experiment 1 averaged across all four stimulus conditions. Plotted on the polar spokes are the rates of pecking to each view of the target geon divided by the rate of pecking to the 0° view of the target geon multiplied by 100 to yield a percentage score. The dashed circle is the mean rate of pecking to all views of all four foil geons divided by the mean rate of pecking to the 0° view of the target geon multiplied by 100. An asterisk indicates that responding to the target was significantly greater than responding to the foil ($p < .05$).

The stimulus generalization gradient for final discrimination training can be characterized by looking at the percentage data as a function of the number of steps the testing view is from the S+. For example, the mean performance of the pigeons one step away from the trained view is calculated by averaging the scores for the -72° and 72° rotations of the target geon. The percentage of pecks for rotation of the target geon one 72° step was 65%, and for two 72° steps it was 57%, indicating a progressive decrement in responding for geon views as they rotate farther from the S+.

To confirm these initial observations, we analyzed the peck data using a full factorial repeated-measures ANOVA of condition (arch+, barrel+, brick+, and wedge+) and degrees of rotation from the target (-144° , -72° , 0° , 72° , and 144°) nested within geon class (target vs. foils). The pigeons showed significantly greater responding to the target geon than to the foil geons as indicated by a statistically significant main effect of geon class, $F(1, 13) = 43.51$, $p < .0001$. A statistically significant main effect of degrees of rotation nested under geon class was also found, $F(4, 64) = 29.99$, $p < .0001$. This finding indicates that responding was affected by the particular viewpoint of the target geon; that is, there was a reliable stimulus generalization gradient. The main effect of condition was not statistically significant, $F(3, 12) = 1.68$, $p > .05$, which indicates that, when collapsed across views, there was no significant difference in responding among the individual conditions. However, the two-way interaction of condition by degrees of rotation nested under geon class was significant, $F(12, 60) = 7.79$, $p < .0001$, indicating that there were differences in responding to specific views of the target geon among the four conditions. Finally, the two-way interaction of condition by geon class, $F(3, 12) = 0.82$, $p > .05$, was not significant, indicating that the difference in responding to the target and foil geons did

not vary among conditions. Post hoc analyses indicated that overall responding to each view of the target geon was significantly above responding to the foils (Newman-Keuls, $\alpha = .05$ for this and later post hoc analyses).

Although the main effect of degrees of rotation was significant, it is possible that this effect was due to a significant drop in responding from the reinforced target view to the nonreinforced views of the target geon; beyond this drop, the pigeons may not have exhibited any differential responding to the nonreinforced views. To test this possibility, data from just the nonreinforced views of the target geon (data from the 0° view and the foils were not included) were analyzed using a full factorial repeated-measures ANOVA of condition (arch+, barrel+, brick+, and wedge+) and degrees of rotation from the target (72° , and 144°). The data were collapsed across the 72° and -72° rotations and across the 144° and -144° rotations to allow for more statistical power. Once again, we found a significant main effect of degrees of rotation, $F(1, 25) = 5.34$, $p < .05$. The interaction between condition and degrees of rotation was also significant, $F(3, 20) = 10.55$, $p < .001$. All other effects were nonsignificant. Post hoc analyses indicated that for the four different conditions, there was a significant decrease in responding between the 72° and 144° rotations in the barrel+, brick+, and wedge+ conditions; there was no significant decrease between the 72° and 144° in the arch+ condition. This analysis indicates that there was a significant decrease in responding as the object was rotated away from the training view to the 72° and 144° views and that the decrease was not limited to a drop in responding between the 0° view and all other views of the target.

Although across all four conditions the pigeons responded significantly more to any and all views of the target geon than they did to all views of the other foil geons, the significant interaction of condition by degrees of rotation indicated that generalization to other views of the target geon differed among the four conditions. Figure 3 shows responding in each of the four individual conditions. The pigeons in the arch+ condition exhibited a steep drop in responding to the nonreinforced views of the target geon (see Figure 3a), indicating no significant generalization to any other views of the target geon above that observed to the foil geons; post hoc analyses indicated that the only view of the target geon that was significantly above the foils was the reinforced 0° view. The pigeons in the barrel+ condition generalized well to the two rotations closest to the reinforced 0° view (see Figure 3b); post hoc analyses indicated that the -72° , 0° , and 72° views were significantly above the foils, but the two farthest rotations, -144° and 144° , were not. The pigeons in the brick+ condition showed relatively consistent generalization across all views (see Figure 3c); post hoc analyses indicated that the -144° , 0° , 72° , and 144° views were significantly above the foils, but the -72° view was not. Finally, the pigeons in the wedge+ condition exhibited a marked asymmetry in generalized responding (see Figure 3d); post hoc analyses indicated that the 0° , 72° , and 144° views were significantly above the foils, but the -72° and -144° views were not.

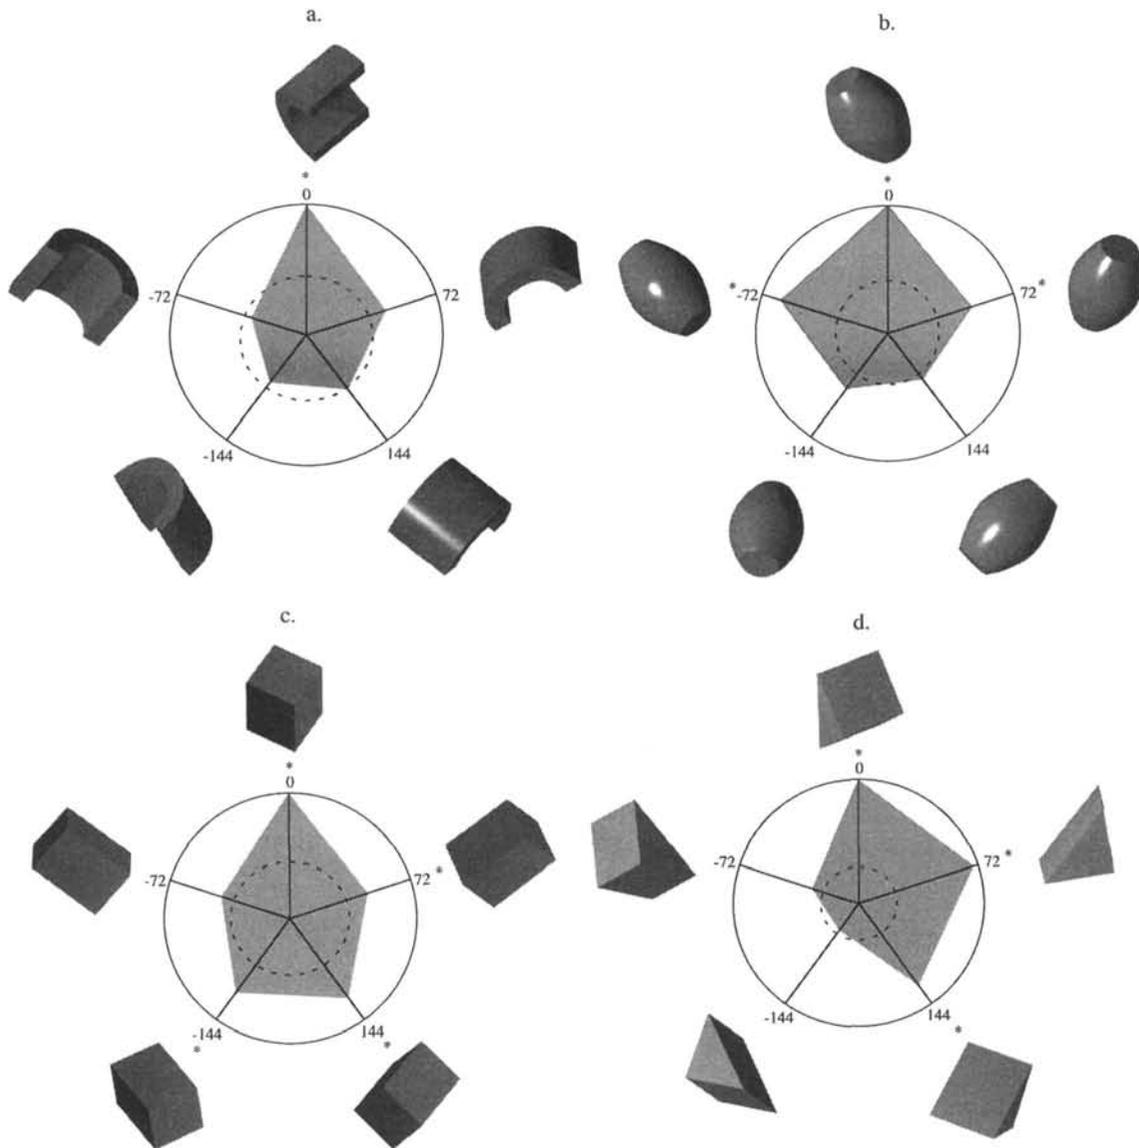


Figure 3. The results for each condition of Experiment 1: arch+, barrel+, brick+, and wedge+. Plotted on the polar spokes are the rates of pecking to each view of the target geon divided by the rates of pecking to the 0° view of the target geon multiplied by 100 to yield a percentage score. The dashed circle is the mean rate of pecking to all views of all four foil geons divided by the mean rate of pecking to the 0° view of the target geon multiplied by 100. An asterisk indicates that responding to the target was significantly greater than responding to the foil ($p < .05$).

Discussion

The data from Experiment 1 indicate that the pigeons did generalize their conditioned responding to other views of a target geon after reinforced training with only one view. On average, the birds responded significantly more to all views of the target geon than to any views of the foil geons, as would be expected by object-centered theories of recognition like RBC.

The pigeons' responding to nonreinforced views of the target geon decreased systematically as the object was rotated farther away from the training view. This systematic

decrease was not simply due to inflated responding at the reinforced view. The failure to obtain viewpoint-invariant recognition is inconsistent with an object-centered theory of object recognition such as RBC. It is important to note, however, that the birds exhibited significant generalization to views as far away as 144° from the target view.

In Experiment 1, the pigeons exhibited different patterns of generalized responding to testing stimuli in the four training conditions. If geons are the basic units of structural object descriptions for pigeons, then they should respond similarly to each stimulus. The differential patterns of responding for the stimuli suggest that the pigeons did not

use these geons as basic units or that the particular task had biased them to respond in a geon-specific manner. In Experiment 2, we sought converging experimental evidence to clarify these outstanding issues.

Experiment 2

Experiment 1 clearly established that pigeons do generalize conditioned responding to other views of an S+ geon, but this generalization was not viewpoint invariant. It must be determined, however, whether the stimulus generalization decrement that we observed to nonreinforced views of the target geon was a result of the particular training procedure that we used in that investigation. In Experiment 1, the birds were reinforced for responding to only a single view of the target geon, and we trained pigeons until they were responding three times as often to that view than to the vast majority (23/24) of the other stimuli. This procedure may have made viewpoint more salient than might normally be the case. To see whether the rotation generalization gradient that we found in Experiment 1 is due to the particular discrimination task or to a more general visual recognition process, we sought converging evidence by using an altogether different discrimination task in Experiment 2.

In Experiment 2, we used a subset of Experiment 1's stimuli in a different discrimination task. We chose a four-alternative forced-choice discrimination task that was more analogous to the naming tasks used to study object recognition in humans and animals. In the task we used in Experiment 2, the pigeons were trained to peck a specific key for a single view of a particular geon. For example, for a particular pigeon, the upper left key was the correct response for the 0° view of the arch, the upper right key was the correct response for the 0° view of the barrel, the lower left key was the correct response for the 0° view of the brick, and the lower right key was the correct response for the 0° view of the wedge (the specific key assignments were counterbalanced across pigeons). Once they became accurate in discrimination training, the pigeons were tested with novel views of each of the four training geons. If, for example, they systematically chose the same key for the 144°, 72°, -72°, and -144° views of the arch as they previously chose for the 0° view, then the pigeons were judged to have generalized the learned response to novel views of the arch; this result would be expected if the pigeons had formed a generic object-centered representation of the object.

On the testing trials with novel geon views, the pigeons were given food reinforcement for any choice they made. Here, unlike Experiment 1, it was neither advantageous nor disadvantageous for birds to respond to the correct key for the untrained views. A decremental rotation generalization gradient under these circumstances would importantly corroborate the gradient that we obtained in the first experiment.

Method

Subjects. The subjects were 4 different feral pigeons (*Columba livia*). By controlled daily feeding, the pigeons were maintained at 85% of their free-feeding weights. Prior to the start of the present experiment, the birds had participated in unrelated studies.

Apparatus. The pigeons were trained in four specially constructed plywood chambers slightly different from those used in Experiment 1. One side of each chamber consisted of a large opening with an aluminum frame attached to the outside of the box. Inside the frame was a clear touch screen (Accutouch Model #002744-FTM-K1; Elographics, Oak Ridge, TN) that was coated with mylar for durability. The pigeon's pecks on the touch screen were processed by a serial controller board (Model #E271-2210, Elographics). A brushed aluminum panel was placed directly in front of the screen to allow pigeons access to limited portions of the video monitor. There were five openings in the aluminum panel. The center opening was a 7 × 7-cm square opening in which the stimuli appeared. The remaining four openings were circular, 1.9 cm in diameter, and they were located 2.3 cm from each corner of the center display opening. The four corner keys served as report keys. In the rear of the chamber, a clear Plexiglas food container was placed level with a wire mesh floor to prevent pigeons from perching on the food cup. Noyes 45-mg pigeon pellets were delivered through a vinyl tube into the food cup using a rotary pellet dispenser (Model #ENV-203M; MED Associates, Lafayette, IN). During experimental sessions, constant illumination was provided by a houselight mounted on the upper rear wall of the chamber. A digital I/O interface board (National Instruments Model #NB-DIO-24; Austin, TX) controlled the pellet dispenser and the houselight.

Control of peripheral stimuli (via the I/O interface) and recording of the pigeons' responses (via the serial controller board) were accomplished by four Apple Macintosh 7100/66 Power PC computers. The pigeon's monitor and an identical monitor located in an adjacent room were connected by a distribution amplifier (Model #MAC/2 DA2; Extron Electronic, Sante Fe Springs, CA). The programs were developed in Hypercard version 2.3.

Stimuli. The stimuli included those used in Experiment 1, with the exception of the cone, which was not used. Additional rotations of the four geons were created at regular 36° intervals.

Procedure. Experiment 2 consisted of three phases: training, Phase 1 testing, and Phase 2 testing. During training, pigeons were taught to peck one of the four corner report keys to the 0° view of each of four geons. For example, Bird 7Y was trained to peck the upper left key for the arch at 0°, the upper right key for the barrel at 0°, the lower left key for the brick at 0°, and the lower right key for the wedge at 0°. Key assignments were counterbalanced across the 4 pigeons. In training, we used a randomized block design; each block consisted of five presentations of each of the four training stimuli for a total of 20 trials in a block. There were 10 blocks in each daily session for a total of 200 trials per day.

At the beginning of a trial, the center display area was illuminated with a black cross centered on a white background. A single peck anywhere within that display area turned on a single-geon stimulus in the center. The pigeons were required to peck the center screen 20 times (observing responses). After the 20th peck, the four corner report keys were illuminated. After a correct choice, the stimulus was removed from the display area, the report keys were darkened, and a food pellet reinforcer was delivered; after an incorrect choice, the stimulus flashed off for 1 s, the report keys were darkened, the houselight was turned off for 4 to 6 s (a mean of 5 s), and one or more correction trials began (these were not scored for analysis). Intertrial intervals ranged from 6 to 10 s (mean of 8 s). Each bird was required to meet a criterion of at least 85% correct overall and at least 80% correct for each of the four geons (an 85/80 criterion) for 3 consecutive days before it could proceed to Phase 1 testing.

In Phase 1 testing, the pigeons saw four novel views of each of the four geons rotated in depth by 72° intervals from the trained 0° view. The trained 0° view for each geon was also given as a testing

Table 2
Percentage Correct in Experiment 2 Phase 1 Testing for Each Bird at Each View Across All Four Geons

Bird	View				
	-144°	-72°	0°	72°	144°
15R	45	58	95	64	54
21B	41	43	94	56	50
35Y	34	53	99	43	34
7Y	50	53	93	68	61
<i>M</i>	43	51	95	58	50

stimulus to allow for comparable data analysis. This testing procedure yielded a total of 20 testing stimuli (-144°, -72°, 0°, 72°, and 144° rotations for each of the four geons). Testing sessions began with eight warm-up trials (two of each training stimulus). The remainder of the testing session followed a randomized block design; each block consisted of 22 trials involving each of the four training stimuli appearing five times and two of the testing stimuli. The two testing stimuli were chosen randomly, without replacement, from the total pool of 20 testing stimuli. All training trials involved differential food reinforcement. Testing trials were nondifferentially reinforced; the pigeons were given reinforcement for any choice on testing trials, so that repeated testing with the same stimuli would not teach the birds the correct response (if they were differentially reinforced) or cause the birds to cease responding to the testing stimuli (if they were not reinforced). There were 10 blocks in each daily session for a total of 228 trials (8 differentially reinforced warm-up trials involving the 0° views, 200 differentially reinforced 0° views, and 20 nondifferentially reinforced testing stimuli). If a pigeon failed to meet the 85/80 criterion for the 208 training trials by the end of a session, then it was returned to training for later sessions until it met the 85/80 criterion. Each pigeon continued until it completed a total of 20 testing sessions of Phase 1 testing. It was then returned to training for 5 days before beginning Phase 2 testing.

In Phase 2 testing, the same 20 testing stimuli from Phase 1 testing were used, and five new views of each geon were added at 36° intervals between the Phase 1 testing stimuli. This addition gave us a total of 40 testing stimuli (-144°, -108°, -72°, -36°, 0°, 36°, 72°, 108°, 144°, and 180° rotations of each geon). Phase 2 testing trials proceeded in much the same manner as Phase 1 testing trials. Testing sessions began with eight warm-up trials (two of each training stimulus). The remainder of the testing session followed a randomized block design; each block consisted of 24 trials involving each of the four training stimuli appearing five times and four of the testing stimuli. The four testing stimuli were chosen randomly, without replacement, from the total pool of 40

testing stimuli. There were 10 blocks in each daily session for a total of 248 trials (8 differentially reinforced warm-up trials involving the 0° views, 200 differentially reinforced 0° views, and 40 nondifferentially reinforced testing stimuli). If a bird failed to meet the 85/80 criterion for the 208 training trials by the end of a session, then it was returned to training for subsequent sessions until it met the 85/80 criterion. The birds continued until they completed a total of 20 testing sessions of Phase 2 testing.

Results

The pigeons took an average of 14.3 days to complete training, 23.0 days to complete Phase 1 testing, and 21.0 days to complete Phase 2 testing. On the last 3 days of training, the pigeons' mean accuracy was 91% correct.

Individual pigeons responded similarly during testing to untrained views of each geon as indicated in Tables 2 and 3. Table 2 gives the percentage of correct choice responses across all four geons for each individual pigeon and mean performance across all 4 pigeons at each view in Phase 1 testing. Table 3 gives the percentage of correct scores across all four geons for each individual pigeon and mean performance across all 4 pigeons at each view in Phase 2 testing. Accuracy generally fell as testing views rotated farther away from the training view. Still, accuracy to the testing views exceeded chance (25%) by from 9% to 33%.

Figure 4 is a polar plot showing the overall accuracy of the pigeons to all four of the geons combined. The results of Phase 1 testing are represented by the black lines connecting the spokes of the graph; where these lines intersect a spoke indicates choice accuracy for that view. For example, in Figure 4 at the 0° view, the line intersects very close to the perimeter of the graph, indicating an accuracy of 95% correct. The results of Phase 2 testing are represented by the gray area of the graph; where the gray area intersects a spoke indicates choice accuracy for that view. The inner circle indicates chance performance (25%) and the dashed circle indicates discrimination performance that is significantly above chance (one-tailed binomial, $p < .05$). As the figure shows, the pigeons were most accurate at the trained 0° view of the geons. Accuracy was lower for the other views of the geons in both phases of testing, but mean performance across pigeons was significantly above chance (25% correct) at all views in both phases of testing (one-tailed binomial, $p < .05$). This result indicates that the pigeons exhibited significant stimulus

Table 3
Percentage Correct in Experiment 2 Phase 2 Testing for Each Bird at Each View Across All Four Geons

Bird	View											
	-180° ^a	-144°	-108°	-72°	-36°	0°	36°	72°	108°	144°	180°	
15R	33	46	60	53	65	90	53	58	43	51	33	
21B	35	36	45	45	49	98	46	53	30	43	35	
35Y	23	38	56	55	51	100	39	41	34	31	23	
7Y	46	78	71	53	65	89	60	68	38	53	46	
<i>M</i>	34	49	58	51	58	94	49	55	36	44	34	

^a-180° = 180°.

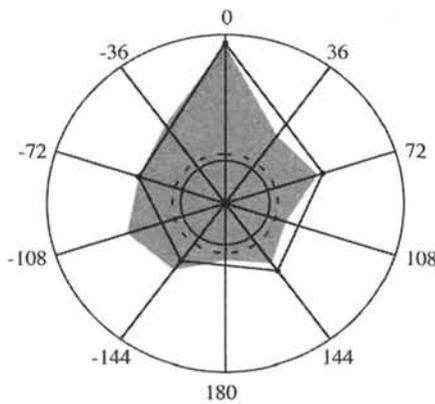


Figure 4. The results of Experiment 2 averaged across geons. Plotted on the polar spokes is the mean choice accuracy (range, 0% to 100%) at each view. The inner circle is chance performance (25% correct) and the dashed circle indicates the statistical threshold for reliably above chance performance (one-tailed binomial test, $p < .05$).

generalization to all other rotations of the geons after training at only a single rotation.

The stimulus generalization gradients for Phase 1 and Phase 2 testing can be further characterized by looking at the testing data as a function of the number of steps the testing stimulus is from the training view. For example, the mean accuracy of the pigeons at one 72° step from the training view is calculated by averaging choice percentages for the -72° and 72° views. For Phase 1 testing, accuracy at the training view was 95% correct across all 4 pigeons. The pigeons' accuracy was 54% correct one 72° step from the 0° view and 46% correct two 72° steps from the 0° view. These data indicate a large generalization decrement from the training view to one 72° step away and a smaller decrease in accuracy as the geon is rotated two 72° steps away from the training view in Phase 1 testing. In Phase 2 testing, accuracy at the training view was 94% correct across all 4 pigeons. One 36° step from the 0° view the pigeons' accuracy was 53% correct, two 36° steps away it was 53% correct, three 36° steps away it was 47% correct, four 36° steps away it was 47% correct, and five 36° steps away it was 34% correct. These data indicate a large decrement was produced by even a very small depth rotation (36° step away) and subsequent decreases in accuracy were smaller as the geon was rotated farther from the trained view, as in Phase 1 testing.

The data from Phase 1 testing were analyzed using a repeated-measures, full factorial ANOVA of geon (arch, barrel, brick, or wedge) by degrees of rotation from the trained view (-144°, -72°, 0°, 72°, and 144°). The pigeons showed a statistically significant main effect for degrees of rotation, $F(4, 12) = 37.23, p < .01$. As indicated by the black lines in Figure 4, this difference confirms a reliable decremental stimulus generalization gradient. The pigeons did not exhibit a significant difference for geon, $F(3, 9) = 1.86, p > .05$. However, a significant interaction between geon and degrees of rotation from the trained view, $F(12,$

36) = 9.86, $p < 0.01$, indicates that stimulus generalization to other views differed across geons.

The data from Phase 2 testing were analyzed using a repeated-measures, full factorial ANOVA of geon (arch, barrel, brick, or wedge) by degrees of rotation from the trained view (-144°, -108°, -72°, -36°, 0°, 36°, 72°, 108°, 144°, and 180°). The pigeons again showed a statistically significant main effect for degrees of rotation, $F(9, 27) = 18.22, p < .01$. As indicated by the gray area in Figure 4, this difference confirms a reliable rotation generalization gradient. The pigeons did not exhibit a significant difference for geon, $F(3, 9) = 1.13, p > .05$. However, a significant interaction between geon and degrees of rotation from the trained view, $F(27, 81) = 4.61, p < .01$, again indicated that stimulus generalization to other views differed across geons.

Figure 5 shows choice accuracy for each individual geon in both Phase 1 testing (filled circles connected by the black line) and Phase 2 testing (gray area). For the arch, the pigeons exhibited modest stimulus generalization to other views (see Figure 5a). Post hoc analyses indicated that the birds responded significantly above chance at the 0° and 144° views in Phase 1 testing and at the -36°, 0°, 36°, 108°, and 144° views in Phase 2 testing (one-tailed binomial, $p < .05$ for this and subsequent post hoc analyses). For the barrel, the pigeons generalized well to most other views, but the generalization gradient was asymmetrical (see Figure 5b). Post hoc analyses indicated that the birds responded significantly above chance at all views with the exception of the 144° view in Phase 1 testing and at the 108° and 144° views in Phase 2 testing. For the brick, the pigeons generalized well to all other views (see Figure 5c). Post hoc analyses indicated that the birds responded significantly above chance at all views in Phase 1 and Phase 2 testing. For the wedge, the pigeons exhibited asymmetrical and irregular generalization (see Figure 5d). Post hoc analyses indicated that the birds responded significantly above chance at the 0°, 72°, and 144° views in Phase 1 testing and at the -108°, -36°, 0°, 72°, and 144° views in Phase 2 testing. Because the individual views of each geon were not related in any systematic way to views of the other geons, the pigeons' incorrect choices were not included in any of the figures (see Appendixes B and C for full tables of the pigeons' errors).

An additional repeated-measures, full factorial ANOVA of phase (Phase 1 testing and Phase 2 testing) by geon by degrees of rotation (only those views given in both phases: -144°, -72°, 0°, 72°, and 144°) was conducted to see if the data from Phase 1 testing differed significantly from the data from Phase 2 testing. The main effect of phase and interactions involving phase were not significant, indicating that the birds did not respond differently from one testing phase to the next.

Discussion

Across all four geons in both testing phases, the pigeons in Experiment 2 exhibited significant rotation generalization after training at only a single view of a geon. Phase 2 testing allowed us to take an even broader look at rotation general-

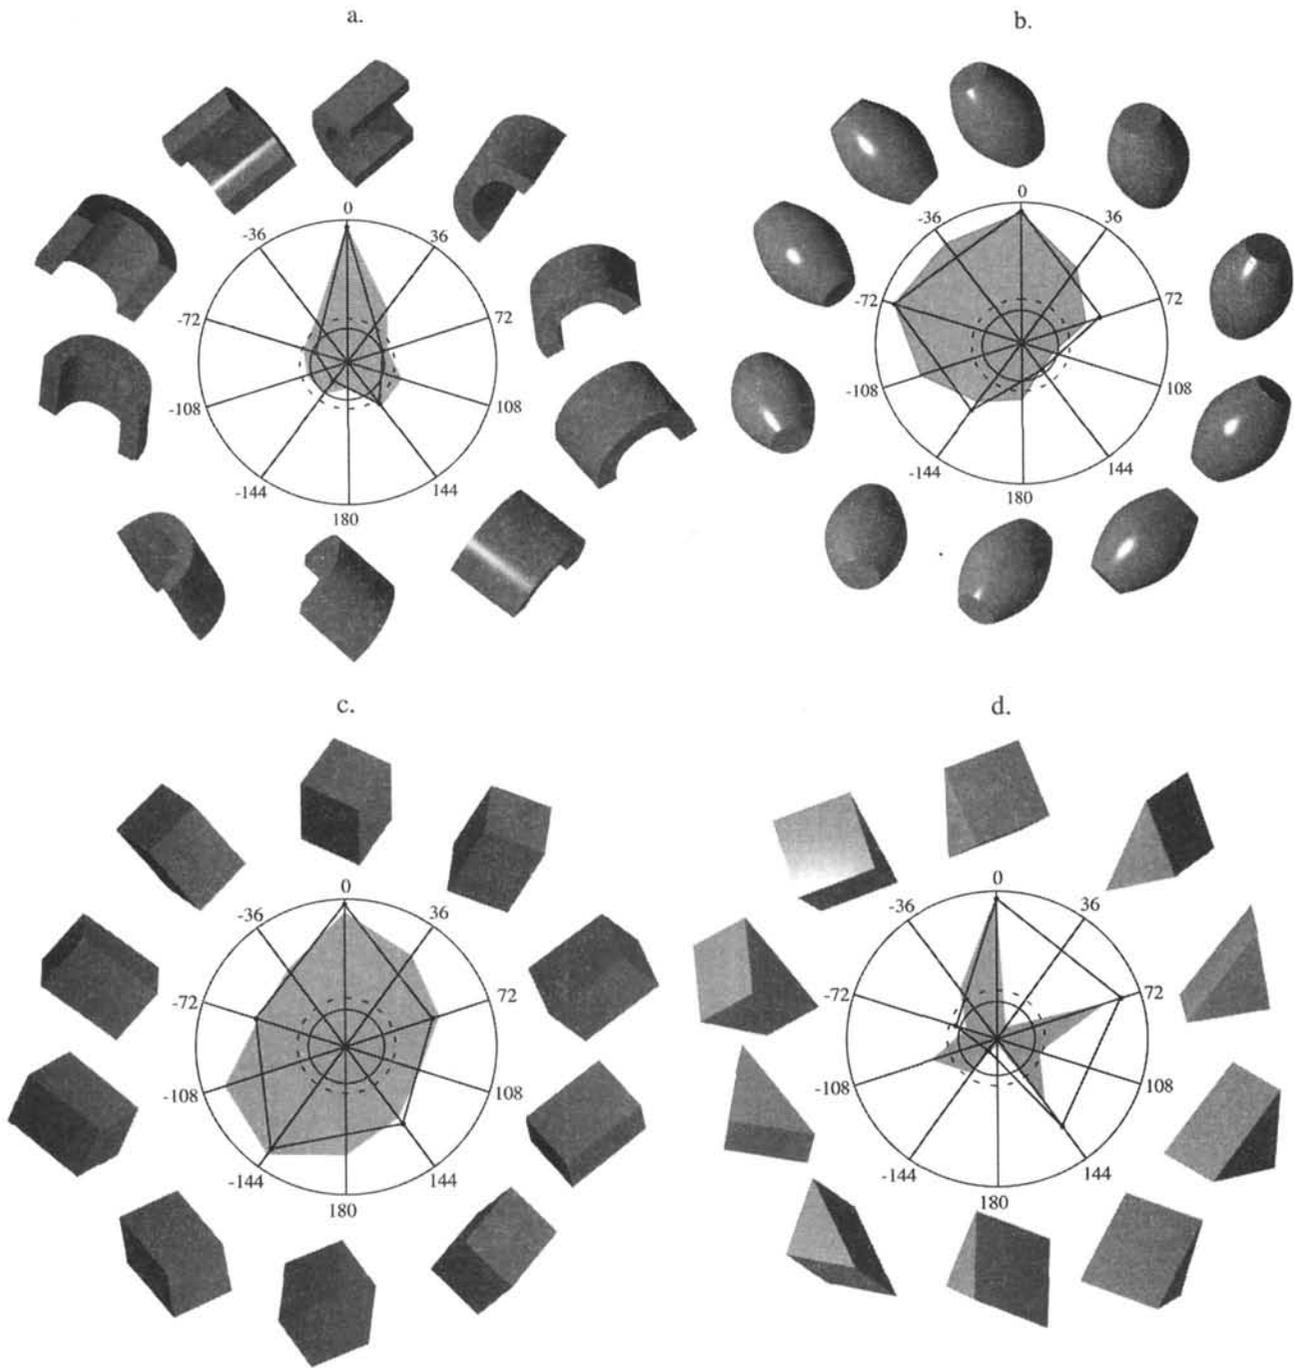


Figure 5. The results for each geon used in Experiment 2. The geons were arch, barrel, brick, and wedge. Plotted on the polar spokes is the mean choice accuracy at each view of each geon. The inner circle is chance performance (25% correct) and the dashed circle indicates the statistical threshold for reliably above chance performance (one-tailed binomial test, $p < .05$).

ization by presenting views that were not used in Experiment 1 or in Phase 1 testing of Experiment 2. Phase 2 testing provided clear evidence of the pigeons' ability to generalize discriminative responding to novel views of the training geons, in some instances to views that were a full 180° from

the training view. Nevertheless, the pigeons also exhibited a systematic decrease in generalization for testing views as they rotated farther away from the trained view. The largest drop in both phases of testing was from the trained view to a view one 36° or 72° step away; subsequent drops were

smaller. Finally, the results from both phases of testing indicate geon-specific responding for different testing views of different geons.

In addition to these results, Experiment 2 answered the main questions that remained from Experiment 1. The pigeons' systematic decrease in responding in Experiment 1 was replicated in Experiment 2, indicating that this result is not due to the particular go/no-go task that we used in Experiment 1. The results of Experiment 2 also confirmed that the geon-specific responding that we found in Experiment 1 was robust and was not task peculiar.

General Discussion

Experiments 1 and 2 of the present study yielded three main results. The first result is that pigeons showed significant stimulus generalization to most other pictured rotations of an object after reinforced training at only one view. The second result is that pigeons exhibited a systematic decrease in stimulus generalization for views of an object as they were rotated away from the training view. The third result is that pigeons showed geon-specific responding, with different geons supporting different stimulus generalization gradients.

Before considering the implications of the data, the generality of these results needs to be considered. First, we discuss the pigeon's ability to process these two-dimensional images as three-dimensional objects. Second, we address the robustness of the results by comparing in greater detail the data from the two different experimental paradigms.

Two-Dimensional Versus Three-Dimensional Processing

Research in other laboratories suggests that pigeons have the ability to perceive two-dimensional stimuli, such as photographs, as representing three-dimensional objects. For example, researchers have found that pigeons can generalize between three-dimensional objects and two-dimensional representations of those objects (Delius, 1992; Lumsden, 1977). Delius (1992) trained pigeons to discriminate between spheres and nonspheres using real objects. In transfer tests, the pigeons exhibited significant generalization to photographs of both familiar and novel objects. These results indicated that the pigeons recognized that the real objects and the photographs were equivalent.

Reid and Spetch (1998) also conducted a set of experiments demonstrating that pigeons perceive three-dimensionality in two-dimensional stimuli. They found that pigeons could use depth cues, such as perspective and shading, to discriminate three-dimensional objects from two-dimensional objects in two-dimensional portrayals of those objects.

On the basis of the results of these experiments, it is plausible to expect that pigeons are capable of perceiving two-dimensional images as representing three-dimensional objects. Although pigeons have the ability to use stimulus cues to perceive three-dimensionality, they may not have done so in our experiments. An alternative explanation of the data in our study is that the pigeons are using simple

two-dimensional features of the images to discriminate among them. These features may include patterns of light and dark or geometric features (e.g., edge location, edge length, or vertex types and their locations). We are unaware of any accepted methods for geometric analysis of the types of figures studied here. Poggio and Edelman (1990) have analyzed paper clip figures in terms of the x, y coordinates of bends in the figure, but this method cannot be extended to describe the rich single-geon objects that we used. We were, however, able to compare the patterns of light and dark in our images through the use of a simple pixel-matching program. Although pixel-matching is an unlikely account of object recognition, it is possible that our pigeons' responding could have been based on an analysis of these low-level features. Therefore, the remainder of our analyses will focus on the explanatory power of a pixel-matching analysis.

To explore pixel-matching as an account of the pigeons' behavior, we analyzed the amount of pixel overlap between the 0° view of each geon with the remaining views of that geon and with all views of the other geons. If a simple pixel-matching analysis is able to account for the generalization that we found in these experiments, then it is possible that the pigeons used the simple two-dimensional attributes of the stimuli rather than attending to the depth cues. If a pixel-matching analysis is unable to account for the pigeon's generalization, however, the pigeons may have applied more sophisticated perceptual processing, such as extracting geometric features or three-dimensional cues in order to perceive depth in the stimuli. We used a modified version of PictCompare to calculate the amount of pixel overlap (we modified PictCompare to quantify the degree of pixel overlap between two Macintosh Pict files). An exact pixel match for a geon with itself yielded a nominal score of 50.

For Experiment 1, the pixel-matching scores were generated by comparing the S- views to the S+ view for each training condition. An overall view of the pixel-matching scores for Experiment 1 is shown in Figure 6. These scores were calculated by averaging the pixel-matching scores across the four target geons, such that the score that resulted

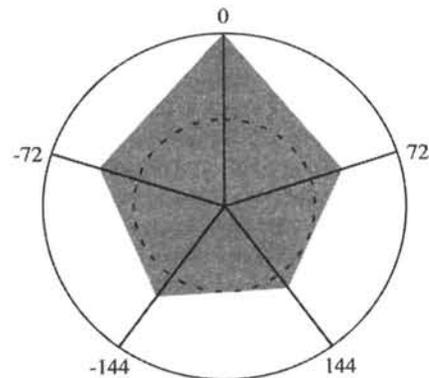


Figure 6. The pixel-matching scores for Experiment 1 averaged across all four conditions. The dashed circle is the mean of the pixel-matching scores for all views of all four foils calculated for each condition and then averaged.

when the -72° view of the arch was compared to the arch at 0° was averaged with the score for the -72° view of the barrel when compared to the barrel at 0° , and so forth. The dashed circle of Figure 6 represents the average of all scores that resulted when the S+ was compared with all of the foils for each condition. For example, the score for the -72° view of the arch compared with the 0° view of the barrel would be included in this calculation. As Figure 6 shows, the pixel-matching scores were generally higher when the S- views of the target geon are compared with the S+ view of the target geon than when views of the different foil geons were compared to the S+ view.

In addition to the scores presented in Figure 6, we also calculated correlations between the pigeons' performance in Experiment 1 and the pixel-matching scores that we computed. We divided the mean rate of pecking to each geon view by the mean rate of pecking to the reinforced view of the target geon and multiplied by 100. These percentages were calculated for each bird individually and then averaged across conditions. A single percentage score for the foils was also calculated for each condition. The mean percentages were then correlated with the corresponding pixel-matching scores; a pixel-matching score for the foils was calculated by taking the mean of the scores for each foil (as compared to the S+ stimulus) in each condition. A Pearson's r was calculated across all four training conditions. Table 4 shows the results of Pearson's r calculations for the pixel-matching scores correlated with the actual pigeon data from Experiment 1. The correlation between the overall percentage of pecks and the overall pixel-matching scores was significant, Pearson's $r(22) = .62$, $p < .01$. Despite the significant correlation, pixel-matching accounted for only 38% of the variance ($R^2 = .38$) in the discriminative performance observed in Experiment 1, leaving 62% of the variance unaccounted for. Comparing Figure 6 (pixel-matching scores) with Figure 2 (actual pigeon performance) suggests a good, but imprecise fit; most obviously, these two functions are oppositely asymmetrical.

Next, we considered how well a pixel-matching analysis accounted for the data in Experiment 2. The pixel-matching scores were now transformed to measure the similarity of each test view to each of the four trained geons. If the test

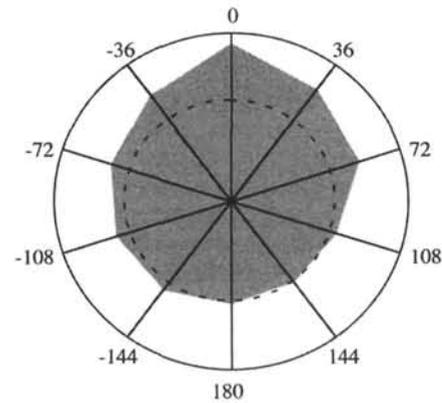


Figure 7. The transformed pixel-matching scores for Experiment 2 averaged across geons. The dashed circle represents the score if pixel-matching yielded equivalent values for all four geons.

stimulus is similar to the noncorresponding geons, then there should be a greater likelihood that the pigeon will make an incorrect than a correct choice response. Thus, the pixel-matching score for the correct choice was divided by the sum of all four pixel-matching scores (one score for each alternative) and then multiplied by 100. The maximum score possible is 100 (50/50 multiplied by 100) and the chance score is 25 (e.g., $x/4x$ multiplied by 100 where x is a pixel-matching score that is identical for all four geons). For example, the 36° view of the arch yielded a pixel-matching score of 34 when compared to the arch at 0° , 25 when compared to the barrel at 0° , 25 when compared to the brick at 0° , and 20 when compared to the wedge at 0° . In this case, we would divide the score for the arch (34) by the total of all four pixel-matching scores (104), multiply by 100, and get a score of 33. These scores were then scaled so that the highest pixel-matching score corresponded with the highest accuracy obtained by the pigeons (e.g., 94% overall in Phase 2 testing).

Figure 7 shows the overall transformed pixel-matching scores. These scores were calculated by taking the mean of the transformed scores for the individual geons. The dashed circle of Figure 7 represents a score of 25, which would be expected if all of the individual pixel-matching scores were equal. At the majority of views, the pixel-matching scores are well above the black line. For the 108° , 144° , and 180° views, however, the pixel-matching scores are only above the line by a small amount. Based on these scores, the birds should not show much, if any, stimulus generalization at these views. The overall bird performance, shown in Figure 4, was above chance at these views, suggesting that the pixel-matching scores underestimated the pigeons' rotation generalization performance at these distant views.

The transformed pixel-matching scores were subsequently correlated with the actual pigeon data from Experiment 2. The pigeons' accuracy for each stimulus was used for this correlation. As the second and third rows of Table 4 indicate, we found a significant overall correlation between the accuracy of the pigeons and the pixel-matching scores in Phase 1 testing, Pearson's $r(18) = .59$, $p < .01$, and Phase 2

Table 4
Correlations of Pixel-Matching Scores With the Pigeons' Performance in Experiment 1 and Experiment 2

Correlation	Geon				
	Overall	Arch	Barrel	Brick	Wedge
Exp. 1 to pixel matching	.62**	.78	.89*	.90*	.42
Exp. 2, Phase 1 to pixel matching	.59**	.78	.95**	.63	.24
Exp. 2, Phase 2 to pixel matching	.57**	.74*	.80**	.67*	.57

Note. Exp. = experiment.
* $p < .05$. ** $p < .01$.

testing, $r(38) = .57, p < .01$. Although the correlations were significant, this pixel-matching analysis accounted for only 35% of the variance in discriminative performance in Phase 1 testing and 32% of the variance in Phase 2 testing, leaving a large percentage of the variance unaccounted for.

The results of the pixel-matching analysis support the idea that, although a simple template-matching technique can account for some of the variance, it does not fully explain the broad generalization exhibited by the pigeons. Thus, the pigeons probably did not rely only on simple features to discriminate among the stimuli and may have used depth cues or other geometric features to extract the three-dimensional nature of the two-dimensional representations.

Generality of Results

Another important aspect of these experiments is the robustness of the reported effects. To determine the robustness of the data, we examined the generality between our discrimination learning tasks. The results of both experiments were remarkably similar, even when comparing pigeon performance for each individual geon across experiments. The response data for each condition in Experiment 1 (Figure 3) and the accuracy data for each geon in Phase 1 testing (Figure 5, filled circles connected by black lines) are nearly identical in shape.

To examine these similarities more precisely, we correlated the results from Experiment 1 with those from Experiment 2. As an index of discriminative performance in Experiment 1, we divided the mean rate of pecking to each geon view by the mean rate of pecking to the reinforced view of the target geon and multiplied by 100.¹ These percentages were calculated for each condition individually to yield a total of 20 numbers (five views of four geons). As an index of discriminative performance in Experiment 2, we used the mean percent correct scores for the 20 stimuli in Phase 1 testing. Only data from Phase 1 testing of Experiment 2 were included because Phase 2 testing included views that were not used in Experiment 1; in addition, data for the cone in Experiment 1 were excluded from the correlation because the cone was not used in Experiment 2.

Overall performance in Experiment 1 was significantly correlated with overall performance in Experiment 2, Pearson's $r(18) = .93, p < .01$, indicating that the results were highly robust across both discrimination learning paradigms. The correlations were also significant for the arch, barrel, and wedge geons considered separately, $r(3) = .95, r(3) = .94, r(3) = .99$, respectively, all $ps < .05$; the correlation was not significant for the brick, but it did approach significance,² $r(3) = .81, p < .10$. The high correlations of discriminative performance between the experiments indicate that the same basic recognition processes were probably at work in both Experiments 1 and 2 despite the very different tasks that each investigation involved.

Result 1: Significant Generalization to Other Views

The results of these two experiments revealed that pigeons generalize recognition responding from one view of a

single geon to most other views of the same geon. This result was anticipated by an object-centered theory, such as RBC (Biederman, 1987). RBC predicts that training with a single view of an object should produce significant generalization to other views of that object as long as the trained view of the object meets three criteria: (a) the object can be decomposed into a GSD, (b) the GSD is distinctive from those of other objects, and (c) the GSD remains stable at the other views of the object (Biederman & Gerhardstein, 1993). Across all four geons, the pigeons exhibited statistically significant stimulus generalization to all tested views in both experiments.

Result 2: Systematic Decrement in Responding

The generalization reported in our study was not viewpoint invariant, as would be predicted by RBC theory; the pigeons showed systematic decreases in recognition responding for other views as the object was rotated away from the training view.

In experiments with human participants, Biederman and Gerhardstein (1993; Experiment 4) have reported viewpoint invariance with single geons rotated up to 90°. We did not see similar viewpoint invariance to novel views of single geons, even when the objects were rotated only 36° from the trained view. The pigeons may not have shown complete generalization because, in both experiments, they were highly overtrained with a single view; overtraining was required for the pigeons to master both tasks. However, our data are consistent with human experiments reported by Tarr et al. (1998), which showed costs for recognition of single geon objects rotated in depth by as little as 45° using three different tasks.

A decrement in recognition due to depth rotation is predicted by viewer-centered object recognition theories (Edelman & Bülthoff, 1992; Tarr & Bülthoff, 1995; Tarr & Pinker, 1989; Ullman, 1989). In viewer-centered theories, the representations of objects that are stored in memory retain much of the viewpoint-specific information that is available when that object is actually encountered. The representations are translation invariant and size invariant, but they must be transformed when that object is encountered at a viewpoint different from the one stored in memory. The farther a rotation is from the stored representation, the more extensive the normalization must be; thus, these theories predict recognition costs as an object is rotated in depth. Some specific viewer-centered theories, such as Tarr and Pinker's multiple-views theory (1989), propose that more than one representation can be stored in memory for frequently encountered objects. A viewer-centered theory might also account for Result 1—significant generalization to novel views—if we suppose that normalization would be adequate for recognition of even very distant rotations. While this is certainly a possibility, it is a post hoc

¹ This calculation was done for each bird individually before overall means were taken.

² The low degrees of freedom in this correlation may be responsible for its not reaching statistical significance.

explanation. Further, it fails to predict precisely when reliable generalization to distant rotations will occur.

While viewer-centered theories readily explain the cause of the systematic decrement in recognition performance, it is more difficult to explain this result in terms of an object-centered theory of recognition, such as RBC. In creating our experimental stimuli, we took great care to avoid accidental views, and we chose the lowest level of object decomposition proposed by RBC—the geon. The pigeons' behavior may simply indicate that they noticed when a geon's rotation differed from the 0° view; they may nevertheless have recognized that the different rotations originated from the same object. In support of this idea, previous research has shown that pigeons are capable of viewpoint-invariant responding when they are trained with multiple views of an object (Peissig, Young, Wasserman, & Biederman, 1999). With multiple-view training, the pigeons are explicitly taught to ignore viewpoint and, thus, appear to be insensitive to viewpoint changes. A viewer-centered theory, however, would also predict viewpoint-invariant responding under these training conditions; additional training views minimize the amount of normalization required for novel views. As is evident from this example, experimentally differentiating between these two theories is problematic because they often make similar predictions.

Result 3: Geon-Specific Responding

In addition to exhibiting significant generalization to other views of training geons and showing a systematic decrease in recognition for other views as they rotated farther from the training view, the pigeons also responded quite differently to different views of each geon. This geon-specificity was evidenced when generalized responding was asymmetrical or much higher or lower for a particular view of a geon. For example, Figure 5d shows the pigeons' accuracy for the wedge in Experiment 2. In Phase 1 testing, the pigeons exhibited above chance responding to the 72° and 144° rotations of the wedge, but not to the -72° and -144° rotations. In Phase 2 testing, we saw this same general pattern, but in addition, the pigeons exhibited poor performance for the 36°, 108°, and 180° rotations of the wedge and above chance performance for the -108°, and -36° rotations. In comparison, Figure 5b shows the pigeons' accuracy for the barrel in Experiment 2. In Phase 1 testing, the pigeons exhibited above chance responding to all views except the 144° rotation. In Phase 2 testing, the pigeons' accuracy was above chance for all views except the 144° and 108° rotations. These same general patterns were also found for the wedge and the barrel in Experiment 1 (Figures 3d and 3b, respectively). Thus, the birds showed different patterns of responding for each geon that are consistent across both experiments.

Neither of the two kinds of theories of object recognition can easily account for these data. RBC would predict this type of object-specific responding for objects which reveal and conceal parts as they are rotated; however, these objects were specifically chosen because Biederman (1987) proposed that these simple structures are the basic units that

make up more complex objects. The definition of a geon is a component which is "readily detectable and relatively independent of viewpoint" (Biederman, 1987, p. 118). It is possible that these specific components do not represent the most basic perceptual units for the pigeons or that the particular demands of the two tasks forced the birds to rely on view-specific mechanisms, such as those proposed by Tarr and Pinker (1989).

The multiple-views model proposed by Tarr and Pinker (1989) predicts systematic generalization decrements to novel views, but not necessarily the type of geon-specific responding that we found in these experiments. It is especially difficult to explain this geon-specific behavior in terms of theories in which mental rotation is the normalization process (Rock & DiVita, 1987; Tarr & Pinker, 1989). Using mental rotation, the distance from the trained view should be the most important factor influencing the pigeons' accuracy for recognition of new views. In some cases, however, a rotation that was farther from the training view was easier for the pigeons to recognize than a rotation that was closer to the training view. For example, in Experiment 2, the pigeons showed higher accuracy to the barrel at -144° than to the barrel at 108° (Figure 5b).

Other, more complex analyses may provide a plausible means for a viewer-centered theory to explain the details of generalization performance. For example, an interpolation method of normalization, such as that proposed by Edelman and Bülhoff (1992), can account for geon-specific responding in cases where there are two or more stored views; the trained views can be interpolated to generate representations that were never actually encountered in training. This normalization process would produce errors based on the accuracy of the interpolated view, which in some cases could lead to the types of object-specific errors that we found. But, our pigeons were trained with only one view.

A variation of the interpolation method may nevertheless be able to account for the data. Vetter, Hurlbert, and Poggio (1995) proposed that an organism may be able to generate "virtual views" based on properties common to a class of objects, such as symmetry. With simple geometric objects like the single geons used in these experiments, one view may contain enough information to generate several novel views. Some novel views may be easier to generate using these rules than other views would be, leading to geon-specific responding. Just what rules should be used to generate "virtual views" of the present experimental stimuli, however, is not at all clear at this time.

Conclusions

The three main results that we documented in this study were: (a) significant rotation generalization to most other views of a trained geon, (b) a systematic decrement in recognition responding to progressive rotations in depth, and (c) geon-specific responding to depth-rotated stimuli. We can conclude from the pigeons' performance in both of our experiments that they used a common recognition process in both tasks and that an object-centered theory is best suited to

explain Result 1, whereas a viewer-centered theory of recognition is best suited to explain Result 2.

Despite the current findings, the breadth of research exploring the pigeon's visual recognition processes has shown that an object-centered approach frequently provides a more complete explanation of the experimental data than a viewer-centered approach. For example, Van Hamme, Wasserman, and Biederman (1992) trained pigeons on a four-key forced-choice task with four line-drawings of objects, in which every other edge and vertex was deleted. The pigeons were then tested with the complementary image (when superimposed, the training image and the complementary image of an object form a complete picture with no overlap in contour). The pigeons responded significantly above chance to the complementary images, even though the images shared no corresponding points or features. These data are consistent with an object-centered theory of recognition, like RBC, because the type of deletion employed in this experiment left enough contour in each portrayal to allow geon recovery. These data are not consistent with a viewer-centered theory that requires an alignment process, however, because there are no corresponding points to use to align a stored representation of the training stimulus with the complementary testing stimulus.

As another example, Wasserman, Kirkpatrick-Steger, Van Hamme, and Biederman (1993) explored the effects of geon configuration on the pigeon's visual recognition of line drawings. Pigeons were trained on a four-key forced-choice task with four line-drawings of objects, each composed of four geons. The pigeons were then tested with stimuli composed of the same four geons, but in new spatial arrangements. These test stimuli retained the height and width of the original training stimuli, as well as the orientation of the parts. The pigeons exhibited a significant drop in accuracy to the scrambled stimuli, from 81% correct for the training stimuli to 53% correct for the testing stimuli. Thus, spatial organization plays a critical role for the pigeons' recognition of complex stimuli, as RBC would predict. Precise predictions from the viewer-centered accounts are more difficult to derive because they have not explicitly examined this experimental manipulation.

Theoretical considerations aside, our chief goal in all of this research is to explore the species generality of object recognition. All species must solve the same problem: recognizing an object when it is seen from varying distances and from many different views, as well as when it is partially occluded by another object. Species that can perform these tasks accurately and efficiently have a greater likelihood of survival than those that cannot. To determine the generality of object recognition processes, data from our experiments can be compared with those in the literature for human participants as well as for other nonhuman species.

The human literature has shown ample evidence for the involvement of both object-centered and viewer-centered mechanisms of object recognition. Biederman (1987) and Biederman and Gerhardstein (1993) reported evidence of object-centered mechanisms of recognition in studies where humans exhibited complete generalization, regardless of view. Rock and DiVita (1987), Tarr and Pinker (1989), Tarr (1995), and Tarr and Bülthoff (1995) reported experiments

in which humans exhibited view-specific performance for untrained views of objects; these data constitute evidence for viewer-centered mechanisms of object recognition. Researchers are now attempting to determine the conditions under which these two mechanisms may be used, with the assumption that the nervous system actually employs two different systems in the recognition process (Jolicoeur, 1990; Stankiewicz, Hummel, & Cooper, 1998).

Similarly, we would like to explore the conditions under which pigeons exhibit behaviors consistent with the use of two different types of recognition mechanisms. Whether or not the pigeon solves object recognition tasks through homologous or analogous processes, discovering ways in which the pigeon visual system operates may provide fresh insights into the human visual system and aid in the development of artificial systems that perform object recognition.

References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.
- Cerella, J. (1977). Absence of perspective processing in the pigeon. *Pattern Recognition*, *9*, 65–68.
- Cook, R. G., & Katz, J. S. (1999). Dynamic object perception by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 194–210.
- Corballis, M. C. (1988). Recognition of disoriented shapes. *Psychological Review*, *95*, 115–123.
- Delius, J. D. (1992). Categorical discrimination of objects and pictures by pigeons. *Animal Learning and Behavior*, *20*, 301–311.
- Edelman, S., & Bülthoff, H. H. (1992). Orientation dependence in the recognition of familiar and novel view of three-dimensional objects. *Vision Research*, *32*, 2385–2400.
- Gross, C. G., Bender, D. G., & Rocha-Miranda, C. E. (1972). Visual properties of neurons in inferotemporal cortex. *Journal of Neurophysiology*, *35*, 96–111.
- Hayward, W. G., & Tarr, M. J. (1997). Testing conditions for viewpoint invariance in object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 1511–1521.
- Jolicoeur, P. (1990). Identification of disoriented objects: a dual-systems theory. *Mind and Language*, *5*, 386–410.
- Kirkpatrick-Steger, K., & Wasserman, E. A. (1996). The what and where of the pigeon's processing of complex visual stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 60–67.
- Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, *5*, 552–563.
- Lumsden, E. A. (1977). Generalizations of an operant response to photographs and drawings/silhouettes of a three-dimensional object at various viewpoints. *Bulletin of the Psychonomic Society*, *10*, 405–407.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London*, *200*, 269–294.
- Peissig, J. J., Young, M. E., Wasserman, E. A., & Biederman, I.

- (1999). The pigeon's perception of depth-rotated shapes. *Cahiers de Psychologie Cognitive*, 18, 657-690.
- Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, 18, 263-266.
- Reid, S. L., & Spetch, M. L. (1998). Perception of pictorial depth cues by pigeons. *Psychonomic Bulletin and Review*, 5, 698-704.
- Rock, I., & DiVita, J. (1987). A case of viewer-centered perception. *Cognitive Psychology*, 19, 280-293.
- Stankiewicz, B. J., Hummel, J. E., & Cooper, E. E. (1998). The role of attention in priming for left-right reflections of object images: Evidence for a dual representation of object shape. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 732-744.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-139.
- 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.
- Tarr, M. J., & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1494-1505.
- Tarr, M. J., Bülthoff, H. H., Zabinski, M., & Blanz, V. (1997). To what extent do unique parts influence recognition across changes in viewpoint? *Psychological Science*, 8, 282-289.
- Tarr, M. J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, 21, 233-282.
- Tarr, M. J., Williams, P., Hayward, W. G., & Gauthier, I. (1998). Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience*, 1, 275-277.
- Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193-253.
- Vetter, T., Hurlbert, A., & Poggio, T. (1995). View-based models of 3D object recognition: Invariance to imaging transformations. *Cerebral Cortex*, 3, 261-269.
- Van Hamme, L. J., Wasserman, E. A., & Biederman, I. (1992). Discrimination of contour-deleted images by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 18, 387-399.
- Wasserman, E. A., Gagliardi, J. L., Cook, B. R., Kirkpatrick-Steger, K., Astley, S. L., & Biederman, I. (1996). The pigeon's recognition of drawings of depth-rotated objects. *Journal of Experimental Psychology: Animal Behavior Processes*, 22, 205-221.
- Wasserman, E. A., Kirkpatrick-Steger, K., Van Hamme, L. J., & Biederman, I. (1993). Pigeons are sensitive to the spatial organization of complex visual stimuli. *Psychological Science*, 4, 336-341.

Appendix A

Percentage of Pecks to the S+ in Experiment 1 for Each Bird at Each View by Condition

Condition and geon	View				
	-144°	-72°	0°	72°	144°
Arch+					
Arch	45	43	100	62	52
Barrel	57	40	44	60	66
Brick	49	45	60	60	56
Cone	49	41	45	42	42
Wedge	35	38	41	41	41
Barrel+					
Arch	48	41	40	45	41
Barrel	53	85	100	66	44
Brick	44	42	40	39	35
Cone	44	35	38	39	38
Wedge	39	42	39	40	39
Brick+					
Arch	48	46	66	45	45
Barrel	47	41	45	45	46
Brick	74	58	100	65	79
Cone	58	44	49	40	53
Wedge	41	40	41	45	44
Wedge+					
Arch	24	22	26	24	24
Barrel	31	25	23	31	36
Brick	28	25	29	27	36
Cone	46	33	54	38	28
Wedge	27	40	100	98	81

(Appendixes continue)

Appendix B

Percentage of Choices Made to Each Key for Each
Test View of All Four Geons in Experiment 2
Phase 1 Testing

Test geon and geon key choice	Test view				
	-144°	-72°	0°	72°	144°
Arch					
Arch	16	23	95	25	36
Barrel	18	39	1	34	26
Brick	64	38	4	39	30
Wedge	3	1	0	3	8
Barrel					
Arch	21	0	0	19	41
Barrel	59	91	94	56	24
Brick	14	4	5	20	16
Wedge	6	5	1	5	19
Brick					
Arch	5	21	3	25	23
Barrel	6	15	1	8	3
Brick	85	63	96	61	65
Wedge	4	1	0	6	10
Wedge					
Arch	15	19	1	8	13
Barrel	48	24	4	3	8
Brick	28	29	0	3	6
Wedge	10	29	95	88	74

Note. Because of rounding, choice percentages may not always add to 100.

Appendix C

Percentage of Choices Made to Each Key for Each Test View of All Four Geons in Experiment 2 Phase 1 Testing

Test geon and geon key choice	Test view									
	-144°	-108°	-72°	-36°	0°	36°	72°	108°	144°	180°
Arch										
Arch	28	29	31	36	96	46	29	39	39	19
Barrel	13	8	14	6	1	4	16	10	20	50
Brick	58	61	43	48	1	50	45	45	39	31
Wedge	3	3	13	10	1	0	10	6	3	0
Barrel										
Arch	26	9	4	6	3	16	30	30	38	26
Barrel	51	73	91	88	93	65	46	23	23	40
Brick	19	11	1	5	3	18	18	35	29	29
Wedge	4	8	4	1	3	1	6	13	11	5
Brick										
Arch	5	9	20	18	3	19	21	18	24	24
Barrel	3	3	14	5	0	1	3	1	1	1
Brick	90	85	61	65	91	76	66	53	60	74
Wedge	3	4	5	13	6	4	10	29	15	1
Wedge										
Arch	24	18	25	28	1	26	15	25	23	20
Barrel	13	1	8	0	1	10	0	5	3	9
Brick	35	35	46	31	1	54	8	40	18	68
Wedge	29	46	21	41	96	10	78	30	58	4

Note. Because of rounding, choice percentages may not always add to 100.

Received October 16, 1998
Revision received September 7, 1999
Accepted September 8, 1999 ■