

Recognizing rotated views of objects: Interpolation versus generalization by humans and pigeons

MARCIA L. SPETCH and ALINDA FRIEDMAN
University of Alberta, Edmonton, Alberta, Canada

Pigeons and humans were trained to discriminate between pictures of three-dimensional objects that differed in global shape. Each pair of objects was shown at two orientations that differed by a depth rotation of 90° during training. Pictures of the objects at novel depth rotations were then tested for recognition. The novel test rotations were 30°, 45°, and 90° from the nearest trained orientation and were either interpolated between the trained orientations or extrapolated outside of the training range. For both pigeons and humans, recognition accuracy and/or speed decreased as a function of distance from the nearest trained orientation. However, humans, but not pigeons, were more accurate in recognizing novel interpolated views than novel extrapolated views. The results suggest that pigeons' recognition was based on independent generalization from each training view, whereas humans showed view-combination processes that resulted in a benefit for novel views interpolated between the training views.

One of the fundamental questions in visual cognition research is how an object is recognized across changes in viewpoint that can drastically alter the visual information supplied to the observer. To survive in a dynamic world, animals and humans alike must be capable of accurately and quickly recognizing objects despite changes in view that occur through movement of the object or of the observer. The existence of these abilities is a given, but the processes that underlie them are less clear, and have been the focus of intensive research and debate for the past couple of decades.

Whether object recognition is accomplished through processes that are viewpoint dependent or viewpoint independent has been intensively investigated in the human literature (see, e.g., Biederman & Gerhardstein, 1993; Edelman & Bülthoff, 1992; Joliceur & Humphrey, 1998; Tarr & Pinker, 1989). According to viewpoint-independent approaches (Biederman, 1987; Biederman & Gerhardstein, 1993, 1995), objects are represented in memory as structural descriptions of the spatial relations among simple, volumetric, three-dimensional (3-D) parts called *geons*. Viewpoint-independent recognition should occur provided that (1) an object's image can be decomposed into geons, (2) the arrangement of the parts provides a structural description that is distinct from that of other arrangements, and (3) the object conserves this structural description across different views. If these conditions are

met, the behavioral signature of viewpoint independence is a recognition function that does not change with distance from the training view.

According to multiple-view theories (Bülthoff & Edelman, 1992; Edelman, 1999; Humphrey & Kahn, 1992; Tarr, 1995; Tarr & Bülthoff, 1995; Tarr & Pinker, 1989), objects are encoded with a viewpoint-dependent frame of reference and are represented as a collection of specific, stored views. Unless the object contains diagnostic features, recognition depends on distance to the nearest represented viewpoint. This prediction is made whether recognition of novel views is assumed to depend on (1) some form of generalization from stored views, (2) a transformation of the novel percept to match the frame of reference (viewpoint) of the nearest stored view (Tarr & Pinker, 1989), or (3) rate of accumulation of activity from neurons selective for the stored object, which should vary systematically as a function of distance between the novel view and stored views (see, e.g., Perrett, Oram, & Ashbridge, 1998). In any case, the behavioral signature of the multiple-view approach is that reaction time (RT) increases and/or accuracy decreases as a function of rotation away from the encoded view(s). The literature on object recognition in humans is extensive, both in data and in theory, but no consensus has yet been reached on whether object recognition is best conceptualized as viewpoint invariant or viewpoint dependent.

The processes that allow animals to recognize rotated objects are also beginning to receive experimental attention (see, e.g., Cook & Katz, 1999; Logothetis, Pauls, Bülthoff, & Poggio, 1994; Peissig, Young, Wasserman, & Biederman, 2000; Spetch, Friedman, & Reid, 2001; Spetch, Kelly, & Reid, 1999; Wasserman et al., 1996). In particular, considerable research has been conducted with pigeons, a visually sophisticated species with an excellent ability to

This research was supported by research grants from the Natural Sciences and Engineering Research Council of Canada. We thank Martha Schultz for assistance with the research. Correspondence and requests for reprints should be addressed to M. L. Spetch, Department of Psychology, University of Alberta, Edmonton, AB, T6E 2E9 Canada (e-mail: mspetch@ualberta.ca).

discriminate, categorize, and remember visual details. In most studies, it has been found that pigeons can recognize objects presented in novel depth rotations, but that their recognition accuracy typically decreases as the object is rotated away from the nearest training view.

Recently (Spetch et al., 2001), we compared object recognition by pigeons and by humans using a design based on work by Tarr, Bühlhoff, Zabinski, and Blanz (1997). Tarr et al. investigated humans' recognition of rotated, paperclip-like objects that either had no distinctive parts or had one, three, or five distinctive parts added. Tarr et al. found that humans showed much stronger viewpoint dependence for zero-part and multipart objects than for one-part objects. We replicated this finding for humans using a simultaneous discrimination procedure, and we tested pigeons with the same stimuli. Recognition functions for pigeons and humans were similar except in two important aspects. First, whereas the humans showed much less viewpoint dependence for the one-part objects than for the zero-part or multipart objects, the pigeons did not show this advantage for one-part objects. Instead, they showed viewpoint dependence for all types of objects. Thus, the pigeons' recognition of an object at novel rotations did not benefit from the presence of distinctive object parts. Other recent work using different objects produced similar results (Spetch et al., 1999).

The second interesting difference between pigeons and humans in the Spetch et al. (2001) study was that the humans, but not the pigeons, showed viewpoint invariance for a novel rotation that was interpolated between the trained views. Both species showed viewpoint dependence outside the training range. The difference between interpolated and extrapolated views for humans replicated the findings of several other researchers (Bühlhoff & Edelman, 1992; Srinivas & Schwoebel, 1998; Tarr, 1995). Better recognition of interpolated than of extrapolated views has also been reported for monkeys (Logothetis et al., 1994) but had not been previously investigated in pigeons. The species-based difference reported by Spetch et al. (2001) suggests the intriguing possibility that pigeons lack view-combination processes that facilitate recognition of interpolated novel views for humans and, possibly, other primates. However, the comparison between interpolated and extrapolated views was based on a single pair of orientations and was confounded with degree of rotation (the interpolated rotation was 30° from the nearest training view, whereas the extrapolated rotation was 45° from a training view). Thus, what appears to be a species-based difference in interpolation processes could instead be a threshold difference.

The present study provides a more systematic investigation of interpolation processes in pigeons and humans. We trained pigeons and humans to discriminate between a pair of shaded 3-D objects presented on a computer monitor. The objects were shown at two depth rotations during training and at eight novel rotations during testing. The novel rotations varied in distance from the nearest training view and were either interpolated between train-

ing views or extrapolated beyond the views experienced during training.

METHOD

Subjects

Ten adult pigeons with varied experimental histories served as subjects. None had previously served in a task involving rotated objects. All the birds were maintained at approximately 85% of their free-feeding weights and housed individually under a 12-h light:dark cycle. Food was obtained during and after experimental sessions. Water and grit were always available in their home cages.

The human subjects were 16 undergraduate students (6 men and 10 women, ranging from 18 to 41 years of age) who received extra credit in introductory psychology for participating in the experiment. Twelve additional subjects were tested but were not included in the reported data. Nine of them were excluded because they failed to meet our accuracy criterion, which was an average score of 75% or better on training trials of Blocks 3 and 4. None of these 9 subjects had achieved accuracy scores above 65% during Training Blocks 1 and 2, and most of them were close to chance level (50%) on all training and test trials. Their performance on test trials was not based on an acquired discrimination between the shapes, and so could not be meaningfully compared to the performance of the pigeons, all of which attained the accuracy criterion prior to testing. Thus, the pigeons were trained to criterion, whereas the humans were selected to criterion. The remaining 3 excluded subjects met the accuracy criterion but were randomly selected for exclusion in order to equate the number of subjects in each group.

Apparatus

The pigeons were tested in operant chambers equipped with color monitors for stimulus display and touch screens for recording pecking responses. The humans were seated approximately 45 cm from a color monitor in a private room, and they responded on a keyboard. For both species, the apparatus was identical to that described in Spetch et al. (2001).

Stimuli

Each stimulus pair consisted of two 3-part objects created with the Adobe Dimensions program (see Figure 1). The stimuli were identical to those used in Spetch et al. (2001) and were similar to those used by Tarr et al. (1997). Each object contained three qualitatively distinct volumes (geons) inserted between two tube-like, nondistinct, 3-D volumes. The geons differed in their nonaccidental properties (see Biederman, 1987, Figure 7). The parts were linked end to end at varying angles. The two objects in each pair differed from each other in both the arrangement of the geons and the angles of attachment. One object in each stimulus pair was designated as the S+, and the other was designated as the S-.

In each presentation, the S+ and S- objects were shown side by side. The orientation of the S+ was the same as the orientation of the S- in all presentations. The initial orientation of each object was oblique with respect to the frontal parallel (picture) plane, and was arbitrarily assigned as the 0° orientation. Across images, the objects were then rotated around the y-axis to create various depth rotations. It should be noted that when objects that are initially oblique with respect to the picture plane are rotated around the vertical axis, they sometimes appear to be tumbling around a different axis. Half of the subjects (Group A) were trained with orientations of 0° and 90°, and then were tested with eight novel orientations (30°, 45°, 60°, 120°, 135°, 270°, 315°, and 330°). The remaining subjects (Group B) were trained with 90° and 180° orientations and were then tested with eight novel orientations (45°, 60°, 120°, 135°, 150°, 210°, 225°, and 270°). Thus, both groups were tested with three novel views that were interpolated between the training views—one rotated by 45°

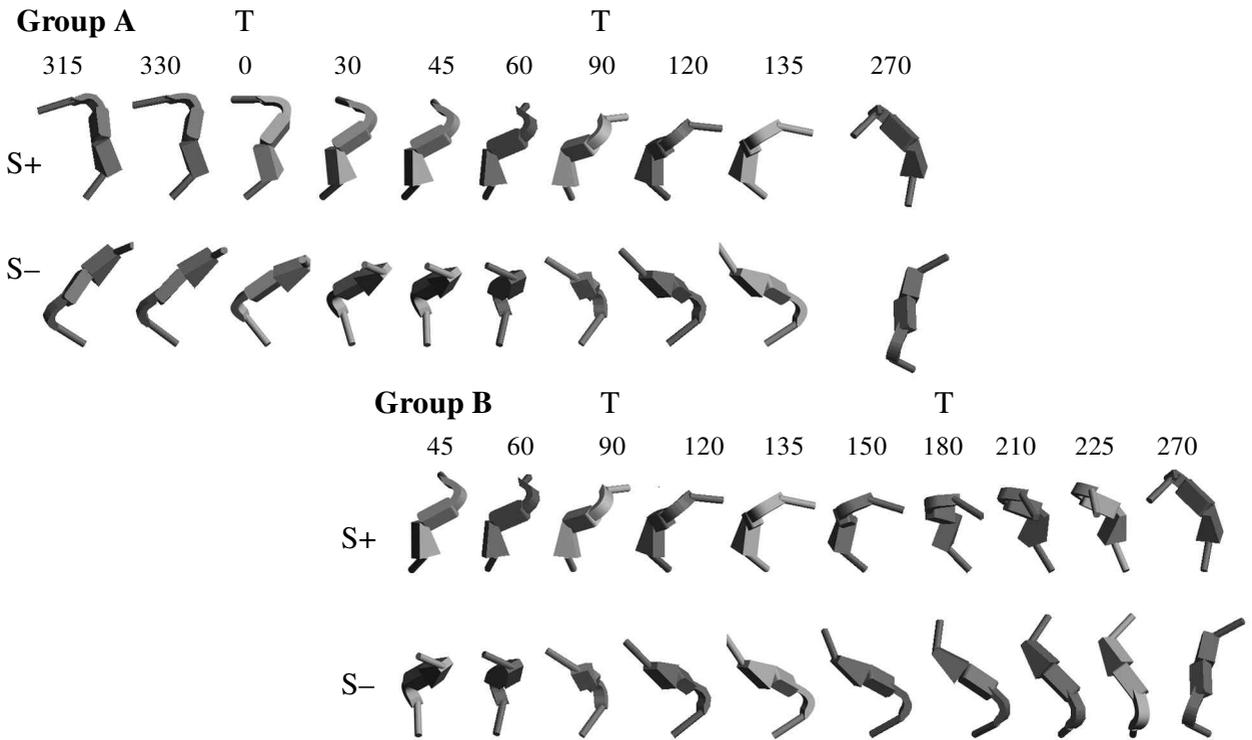


Figure 1. Positive (S+) and negative (S-) objects at each orientation used in training and testing for each group. The numbers indicate the orientation of the objects, in degrees, and T indicates the orientations presented during training.

(in45) and two rotated by 30° (in30) from the nearest training view—as well as five novel views that were extrapolated beyond the training range—two rotated by 30° (ex30), two rotated by 45° (ex45), and one rotated by 90° (ex90) from the nearest training view. However, the specific orientations were counterbalanced across groups.

Procedure

The general procedures were the same as those described in detail in Spetch et al. (2001). For both species, a randomly selected half of the subjects were assigned to Group A, and the remaining subjects were assigned to Group B.

Procedure for pigeons. The task was a simultaneous discrimination task between an S+ and an S- object. Each trial began with the presentation of a stimulus pair that remained visible until the pigeon pecked once at one of the objects. A peck at an object was defined as a peck on the screen that fell within the 4 × 5 cm rectangular area containing the object. A peck at the S+ object was rewarded with food, and a peck at the S- object terminated the trial without food. The objects were presented equally often at each of the two training orientations (0° and 90° for Group A and 90° and 180° for Group B). For each orientation, the S+ was equally often on the left and on the right. To prevent a position bias, incorrect choices were followed by correction trials, in which the same stimulus pair was re-presented until the pigeon made a correct choice. The data from the correction trials were not used in any analyses. The sessions consisted of 120 non-correction trials, plus as many correction trials as were needed. Trials were separated by a 2-sec intertrial interval, during which the monitor was blank. Training continued until the birds met an accuracy criterion of 80% or greater over a block of five consecutive sessions. The birds then received five test sessions, each consisting of 60 unreinforced test trials mixed randomly among 60 reinforced

baseline trials at the training orientations. On test trials, the object pairs were presented equally often at the two training orientations and at the eight novel orientations. At each orientation, the S+ was equally often on the left and on the right.

Procedure for humans. The subjects were seated at the computer and instructed that (1) their task was to choose among pairs of objects using arrow keys on the keyboard, (2) they should use the feedback to learn which object was correct and then respond as quickly and accurately as possible, (3) they would earn points for being correct and double points for being correct and fast, and (4) some trials would end without feedback. The simultaneous discrimination task was similar to that used for the pigeons. The subjects received one point for each correct choice plus one extra point for each correct choice that occurred within 0.5 sec of onset of the stimulus display. The session consisted of four blocks. In Block 1, 8 training trials were presented, all of which were followed by a feedback display (“Correct” or “Wrong,” and cumulative points earned). In Block 2, 8 training trials were presented, half of which were followed by feedback. Trials without feedback ended with the message “No feedback trial.” Blocks 3 and 4 each consisted of 32 trials, 12 of which were training trials with feedback and 20 of which were test trials without feedback. On test trials, the two training orientations and the eight test orientations were each presented twice, once with the S+ on the right and once with the S+ on the left. The order of presentation of trial types within each block was randomly determined for each subject. The accuracy and latency of each response were recorded.

Data Analyses. All of the data presented are from the test trials. An alpha level of .05 was used for all statistical tests. Because the absolute orientations used in training (i.e., 0° and 90° or 90° and 180°) were exactly counterbalanced over subjects for both species, the data were collapsed across groups as a function of rotation from the training views. For example, the first training orientation pre-

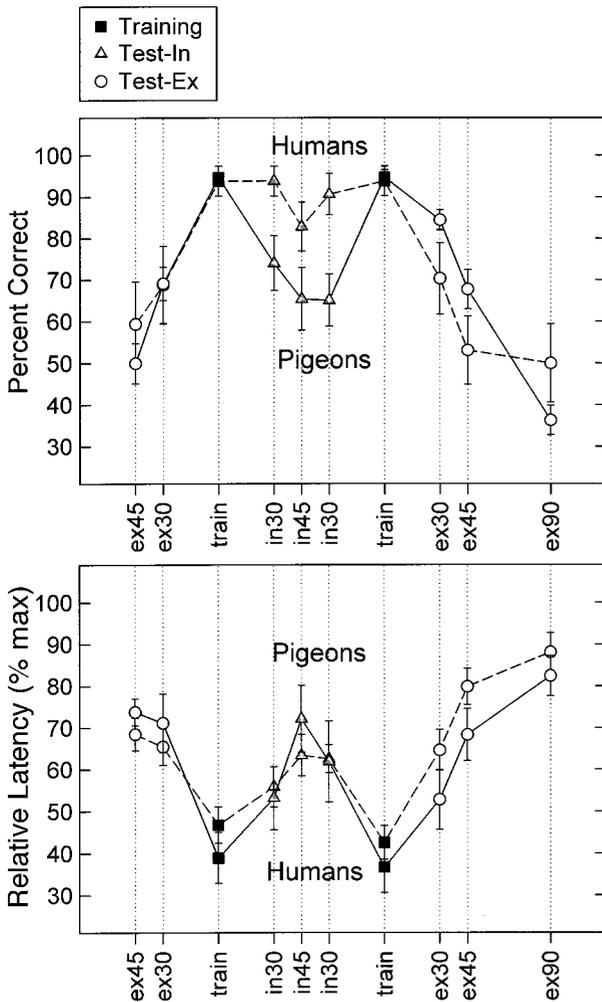


Figure 2. Choice accuracy (top panel) and relative latency (bottom panel) in pigeons (solid lines) and humans (dashed lines) on test trials with each of the trained and novel orientations. Error bars show standard errors of the means. In, interpolated. Ex, extrapolated.

sented in Figure 2 represents the 0° trials for Group 1 and the 90° trials for Group 2, averaged together. The in45 rotation is an average of the 45° trials for Group 1 and the 135° trials for Group 2. For statistical analyses, the data were further collapsed by averaging across novel orientations that provided the same type and degree of rotation. Specifically, for each subject of each species, the interpolated rotation of 30° , the extrapolated rotation of 30° , and the extrapolated rotation of 45° were each represented by two values, which were averaged.

The latency data include both correct and incorrect trials, and, hence, represent the overall time taken by the subjects to choose between the stimuli, whether correctly or incorrectly. The more typical RT measure, which excludes latencies on incorrect trials, was not analyzed, because the human subjects received only a few test trials at each orientation and the exclusion of incorrect trials resulted in missing cells. Some pigeons occasionally had very long latencies that, presumably, reflected off-task activities. Therefore, on a session-by-session basis, we excluded trials on which the latency scores were more than 3 standard deviations above the mean. Because the pigeons had multiple sessions of testing, each one still had latency data from several trials for all trial types. A similar exclusion of outliers

could not be performed for the data from humans, because there were too few trials of each type.

Absolute latencies could not be meaningfully compared across species because of differences in the response measures (pecking vs. pressing the keyboard) and instructions (i.e., the humans, but not the pigeons, were instructed to respond as quickly as possible). Therefore, we scaled the latency data. For each subject, we identified which of the 10 test rotations (i.e., the 2 training rotations and 8 novel rotations) had the highest average latency. We then transformed the latency scores for each test rotation into a percentage of this maximum score and analyzed these relative latencies analogously to the accuracy data.

RESULTS

Accuracy

The primary question of interest was whether recognition of the interpolated versus the extrapolated novel views differed qualitatively for pigeons and for humans. The accuracy scores very clearly indicated a species-basis difference. Specifically, the humans remained very accurate for interpolated rotations but showed large decrements in accuracy for extrapolated rotations, whereas the pigeons showed equivalent decrements in accuracy for interpolated and extrapolated rotations. The top panel of Figure 2 shows the accuracy of each species at the two training views and at each of the novel rotations.

To confirm the apparent species-based difference between interpolated and extrapolated rotations, we conducted a mixed-design analysis of variance (ANOVA), using only the data from the 30° - and 45° -rotation tests. This allowed a factorial design with species as the between-subjects factor and degree (30° vs. 45°) and view type (interpolated vs. extrapolated) as within-subjects factors. This ANOVA revealed significant main effects of degree [$F(1,24) = 48.28$] and view type [$F(1,24) = 4.86$], as well as a significant interaction between species and view type [$F(1,24) = 5.13$]. No other main effects or interactions were significant. The species \times view type interaction confirmed the striking difference between the species: The humans, but not the pigeons, were more accurate with interpolated than with extrapolated views. Both species were less accurate with 45° rotations than with 30° rotations.

Two sets of additional, a priori comparisons provided corroborating evidence of the species-based difference in reaction to extrapolated versus interpolated views. First, relative to training views, the humans did not show a significant drop in accuracy on either of the interpolated view tests [in30, $t(15) = 0.38$; in45, $t(15) = 1.60$], whereas the pigeons showed a significant drop in accuracy on both interpolated tests [in30, $t(9) = 3.56$; in45, $t(9) = 3.42$]. Second, we conducted two-sample t tests to compare accuracy by humans and that by pigeons at each of the rotation conditions (training, in30, ex30, in45, ex45, and ex90). These tests revealed a significant species-based difference only for the in30 tests [$t(24) = 3.57$]. The species-based difference approached significance on the in45 tests [$t(24) = 1.83$], but was not close to significance with training views or on any other tests (all $ps > .25$). Thus, the ac-

curacy of the pigeons was very similar to that of the humans except on interpolated tests.

Latency

The bottom graph in Figure 2 shows the relative latency for each species at each rotation and view type. Latency, like accuracy, showed a robust effect of degree of rotation for both species. However, the difference between the interpolated and the extrapolated views, seen in the humans' accuracy scores, was less pronounced in their latency scores. A mixed-design ANOVA using only the data from the 30°- and 45°-rotation tests revealed a significant main effect of degree only [$F(1,24) = 13.16$].

The a priori comparisons of latency on interpolated views relative to training views revealed that both species showed a significant increase in latency on both interpolated view tests [in30, $t(15) = 3.65$ and in45, $t(15) = 4.10$ for humans; in30, $t(9) = 4.02$ and in45, $t(9) = 4.27$ for pigeons]. Moreover, two-sample t tests comparing the relative latency of the humans and that of the pigeons on each rotation condition failed to reveal any significant species-based differences. Thus, the latency measure was not sensitive to the species-based differences seen in the accuracy measure.

DISCUSSION

In accordance with previous research (Bülthoff & Edelman, 1992; Srinivas & Schwoebel, 1998; Tarr, 1995), the humans showed better recognition of novel views that were interpolated between two training views than of novel views that were extrapolated outside of the training range, even when degree of rotation was equated. Interestingly, this facilitated recognition of interpolated views was robust for the measure of accuracy, but was not significant in the latency measure, which showed viewpoint dependence for both interpolated and extrapolated views. Edelman (1999, p. 172) has suggested that error rate effects may provide a better insight into the workings of a recognition system than RTs do, and are less likely to dissipate with practice.

The pigeons, by contrast, showed strong viewpoint dependence in both accuracy and latency for both interpolated and extrapolated novel views. Moreover, they showed no significant facilitation of recognition for interpolated views over extrapolated views with either measure. Instead, their recognition accuracy and speed appeared to be affected by degree of rotation only.

The benefit shown by humans for interpolated over extrapolated views has been discussed in terms of various conceptions of view-combination mechanisms (see, e.g., Bülthoff & Edelman, 1992; Edelman, 1999). For example, Ullman and Basri (1991) proposed a linear combination of 2-D views; Poggio and Edelman (1990) and Bülthoff and Edelman (1992) proposed a two-stage combination of linear and nonlinear mechanisms operating on 2-D views, which they referred to as *view approximation*; more recently, Edelman (1999) has proposed a variant of the

view-approximation theory in which neural units are tuned to particular views of an object and the basic operation is interpolation between such views. Although the latter is principally a multiple-view approach, it can exhibit viewpoint invariance for novel objects by interpolating between the view spaces of the familiar ones. In Tarr and Pinker's (1989; see also Tarr, 1995) approach, the difference between interpolated and extrapolated views can be accounted for by assuming that subjects mentally rotated between the training views during the training trials.

All of the mechanisms proposed to underlie the benefits for interpolated versus extrapolated views presuppose that the two views to be combined are mapped onto the same representation. On the basis of preliminary data, Spetch et al. (2001) suggested that pigeons may not perceive the two training views as representing the same object. Our comparison of interpolated and extrapolated views provides strong support for that suggestion. With degree of rotation equated and specific orientations counterbalanced, the pigeons did not recognize interpolated novel views more readily than extrapolated novel views, suggesting that they did not benefit from view-combination processes. Rather than mapping the two training views onto the same representation, the pigeons apparently formed two independent representations of the S+ object as it appeared in each training view. During tests with novel orientations, their choices appeared to be determined by generalization from each independent representation.

A process of generalizing from independent representations at each training view would be expected to produce superior recognition of interpolated views only if there is overlap between the two generalization functions. If the representations at the two training orientations overlap, then novel orientations between the views should receive generalization from both views, whereas novel extrapolated orientations could receive generalization from only one view. However, if the training views are far apart in subjective space, or if the generalization functions around each are narrow, then there may be little or no overlap in the functions. In that case, interpolated views would not benefit from generalization to more than one view except, possibly, at the midpoint between the two views. One way of testing a generalization account would be to use training stimuli that differed from each other in increasingly smaller increments, to see whether two views of the same object would eventually produce better performance on an interpolated stimulus than on an extrapolated stimulus.

It is important to note that the pigeons' failure to show view combination does not represent a general inferiority of object recognition or generalization processes. The pigeons learned to discriminate the objects as accurately as the humans did (albeit with more training), and, importantly, they generalized just as well, if not better, to the extrapolated views. This argues against the possibility that the species-based difference reflected a procedural artifact due to differences in reinforcement conditions, correction trials, or duration of training. Such variables might affect general performance levels on test trials but would not

be expected to interact with the comparison between interpolated and extrapolated rotations. Instead, the species-based difference we observed appears to be specifically related to the benefit accrued by view-combination processes.

The finding that the pigeons were as good as or better than the humans in recognizing novel extrapolated views also suggests that the difference between species is qualitative rather than quantitative. Consider the hypothetical generalization functions described above. Differences between species in the occurrence of an interpolation advantage could arise from a quantitative difference in which generalization functions are broad and overlapping for humans but are narrow and nonoverlapping for pigeons. If this were the case, however, one would expect to see less generalization decrement for humans than for pigeons at novel extrapolated orientations, because the broader functions should produce better recognition over a wider range of novel orientations. The fact that the humans showed as much disruption in recognition accuracy as did the pigeons when exposed to extrapolated views argues against the possibility and suggests a qualitative difference in recognition processes.

An interesting question for future research is whether pigeons can be shown to benefit from view combination if procedures are used that may encourage the mapping of training views onto a single representation. For example, training with more closely spaced views, training with dynamically rotating objects, or training with actual 3-D objects may facilitate pigeons' recognition of the correspondence between different views and, hence, may allow them to benefit from view-combination processes.

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 & Performance*, **19**, 1162-1182.
- BIEDERMAN, I., & GERHARDSTEIN, P. C. (1995). Viewpoint-dependent mechanisms in visual object recognition: Reply to Tarr and Bülthoff (1995). *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 1506-1514.
- BÜLTHOFF, H. H., & EDELMAN, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences*, **89**, 60-64.
- COOK, R. G., & KATZ, J. S. (1999). Dynamic object perception by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 194-210.
- EDELMAN, S. (1999). *Representation and recognition in vision*. Cambridge, MA: MIT Press.
- EDELMAN, S., & BÜLTHOFF, H. H. (1992). Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, **32**, 2385-2400.
- HUMPHREY, G. K., & KHAN, S. C. (1992). Recognizing novel views of three-dimensional objects. *Canadian Journal of Psychology*, **46**, 170-190.
- JOLICŒUR, P., & HUMPHREY, G. K. (1998). Perception of rotated two-dimensional and three-dimensional objects and visual shapes. In V. Walsh & J. Kulikowski (Eds.), *Visual constancies: Why things look as they do* (pp. 69-123). Cambridge: Cambridge University Press.
- LOGOTHETIS, N. K., PAULS, J., BÜLTHOFF, H. H., & POGGIO, T. (1994). View-dependent object recognition by monkeys. *Current Biology*, **4**, 401-414.
- PEISSIG, J. J., YOUNG, M. E., WASSERMAN, E. A., & BIEDERMAN, I. (2000). Seeing things from a different angle: The pigeon's recognition of single geons rotated in depth. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 115-132.
- PERRETT, D. I., ORAM, M. W., & ASHBRIDGE, E. (1998). Evidence accumulation in cell populations responsive to faces: An account of generalisation of recognition without mental transformations. *Cognition*, **67**, 111-145.
- POGGIO, T., & EDELMAN, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, **343**, 263-266.
- SPETCH, M. L., FRIEDMAN, A., & REID, S. (2001). The effect of distinctive parts on recognition of depth-rotated objects by pigeons and humans. *Journal of Experimental Psychology: General*, **130**, 218-225.
- SPETCH, M. L., KELLY, D. M., & REID, S. (1999). Recognition of objects and spatial relations in pictures across changes in viewpoint. *Current Psychology of Cognition*, **19**, 729-764.
- SRINIVAS, K., & SCHWOEBEL, J. (1998). Generalization to novel views from view combination. *Memory & Cognition*, **26**, 768-779.
- 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 by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception & 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 Science*, **21**, 233-282.
- ULLMAN, S., & BASRI, R. (1991). Recognition by linear combinations of models. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **13**, 992-1005.
- WASSERMAN, E. A., GAGLIARDI, J. L., COOK, B. R., KIRKPATRICK-STEGE, K., ASTLEY, S. L., & BIEDERMAN, I. (1996). The pigeon's recognition of drawings of depth-rotated stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **22**, 205-221.

(Manuscript received September 26, 2001;
revision accepted for publication February 22, 2002.)