



Shape from shading in pigeons

Robert G. Cook^{*}, Muhammad A.J. Qadri, Art Kieres, Nicholas Commons-Miller

Department of Psychology, Tufts University, USA

ARTICLE INFO

Article history:

Received 8 July 2011

Revised 3 February 2012

Accepted 4 May 2012

Available online 29 June 2012

Keywords:

Shape from shading

Visual perception

Pigeon

Comparative perception

ABSTRACT

Light is the origin of vision. The pattern of shading reflected from object surfaces is one of several optical features that provide fundamental information about shape and surface orientation. To understand how surface and object shading is processed by birds, six pigeons were tested with differentially illuminated convex and concave curved surfaces in five experiments using a go/no-go procedure. We found that pigeons rapidly learned this type of visual discrimination independent of lighting direction, surface coloration and camera perspective. Subsequent experiments varying the pattern of the lighting on these surfaces through changes in camera perspective, surface height, contrast, material specularity, surface shape, light motion, and perspective movement were consistent with the hypothesis that the pigeons were perceiving these illuminated surfaces as three-dimensional surfaces containing curved shapes. The results suggest that the use of relative shading for objects in a visual scene creates highly salient features for shape processing in birds.

© 2012 Published by Elsevier B.V.

1. Introduction

Light is the origin of all vision. While the role of edges and junctions has received the majority of attention in trying to understanding the processes of three-dimensional (3D) object perception and recognition (Hubel & Wiesel, 1962; Marr, 1982), natural objects in the world also typically have continuous and curved surfaces that differentially reflect light. The resulting patterns of shading, shadow, and specularity all contain highly useful shape information for any complex nervous system attempting to internally represent its surrounding visual environment (Norman, Todd, & Orban, 2004). These types of lighting information make essential contributions to the human perception of objects, depth, and the spatial layouts of scenes (Gibson, 1950, 1979; Mingolla & Todd, 1986; Ramachandran, 1988), and are involved with algorithms for machine vision (Horn & Brooks, 1989; Zhang, Tsai, Cryer, & Shah, 1999), and graphic visualization (Kniss, Premoze,

Hansen, Shirley, & McPherson, 2003). Despite its importance in human perception, the incorporation of these types of information remains a theoretical challenge for computational theories of vision.

Given its potential functional value, however, the topic of shape-from-shading has not received a great deal of systematic experimental attention in many animals (Imura & Tomonaga, 2003; Tomonaga, 1998). With their heavy dependence on vision for such essential natural behaviors as foraging and navigation, it would seem natural that birds might be keenly sensitive to the relative patterns created by the differential shading, shadows and specularity of light reflected from surfaces and objects. Because of the computational challenges of processing shaded information, examining these issues in birds is especially interesting because of the small absolute size of their central and visual nervous systems. Are birds sensitive to this type of visual information? If so, how is this processed and implemented within their limited nervous system? Would the mechanisms or perceptual biases involved be the same or different from those proposed for human and computer visual processing of this type of information? Here we report the results of five experiments testing pigeons discriminating concave and convex 3D surface direction as produced from differential lighting and shading.

^{*} Corresponding author. Address: Department of Psychology, Tufts University, 490 Boston Ave., Medford, MA 02155, USA. Tel.: +1 617 627 2456; fax: +1 617 627 3181.

E-mail address: Robert.Cook@tufts.edu (R.G. Cook).

Little is known about how lighting contributes to visual perception in birds. Several early attempts focused on how chicks learn to deal with illuminated objects during development. Hess (1950) found that differential developmental experience during feeding with lighting from above or below influenced the subsequent selection of a photograph of grains illuminated from these directions. Hershberger (1970) trained chicks to discriminate concave and convex dents in actual objects. In contrast, developmental conditions in which light consistently came from above or below did not produce a difference when transferred to pictures of these dented surfaces.

More recently, Reid and Spetch (1998) examined how depth from shading and perspective cues contributed to an object discrimination by pigeons. Pigeons were trained to discriminate images of objects in which the depth cues from shading or perspective were either absent or altered. In transfer tests with novel objects, the pigeons performed above chance with perspective-only and shading-only trials, although the latter was not particularly robust. These results suggest that both of these cues may contribute to the pigeons' perception of depth in pictures.

Young, Peissig, Wasserman, and Biederman (2001) investigated how variations in surface illumination influence the discrimination of geon-based objects in pigeons. In a choice task, the pigeons showed sensitivity to the relative pattern of light reflected from different computer-rendered objects. When the illumination source of the objects was rotated about 90° from that experienced during training, the pigeons showed above chance transfer to these new conditions, but simultaneously exhibited a considerable drop in discrimination performance. The latter fact suggests that features created from the specific pattern of lighting on the objects were involved in the representation of the objects (cf. Tarr, Kersten, & Bülthoff, 1998). The generalized role of such surface illumination in the pigeons' representation and perception of these objects may be limited by the fact that only a small set of stimulus exemplars were used in these experiments, making image memorization possible (Cook, Levison, Gillett, & Blaisdell, 2005; Fagot & Cook, 2006).

Testing pigeons, Cavoto and Cook (2006) examined how different monocular depth cues influenced the depth ordering of computer-rendered objects presented on a textured surface designed to simulate a visual scene. Testing with large number of exemplars, they established that occlusion, relative size, and relative density all contributed to the accurate depth-ordering of the objects by these birds. They also found that lighting played a role in the depth processing of the scenes as the differential illumination of the objects' surface also contributed to the pigeons' performance.

Finally, Rowland, Cuthill, Harvey, Speed, and Ruxton (2008) recently used a field approach to investigate the potential role of shading in natural settings. Artificial prey having normal countershading were eaten less often than uniformly-colored prey by free-living avian predators. This suggests that the differential shading on the ventral and dorsal surfaces of many animals may help camouflage them by making them blend more seamlessly into the surrounding illumination.

This heterogeneous group of studies clearly hints that object and scene illumination cues are potentially important and valuable to birds when determining relative depth and object shape. Given the varied approaches, summarizing the visual and cognitive mechanisms involved and their similarity to those in humans, other primates (Imura & Tomonaga, 2003; Tomonaga & Imura, 2010), or computers is difficult. Because of its potential saliency to avian vision, the goal of these new experiments was to more comprehensively examine how pigeons process differential lighting and shading. To accomplish this, we used the approach of testing computer-rendered artificial stimuli (see Fig. 1). While clearly synthetic, the stimuli had characteristics that would be similar to the experiences of a bird in flight by providing an aerial perspective onto a receding surface or landscape-like scene.

Here we report the results of five experiments testing the sensitivity of pigeons to discriminating the surface direction of concave and convex 3D shapes as produced from differential lighting and shading. In Experiment 1, the pigeons were trained to discriminate between these different surface shapes using a go/no-go procedure. We then investigated how various illumination and surface conditions influenced this shape-from-shading discrimination. Experiment 2 examined how changes in camera perspective and relative surface height influenced this discrimination. Experiment 3 examined the effects of different aspects of surface illumination. Experiment 4 examined performance with novel sets of convex and concave surface shapes. Finally, Experiment 5 examined the effects of dynamic and continuous illumination and perspective changes on behavior. Collectively, the experiments suggest that the differential lighting from these surfaces made important contributions to the object, depth, and scene perception by the pigeons.

2. Experiment 1

Experiment 1 tested the pigeons' ability to discriminate between 3D concave and convex surfaces created from differential shading. The convex stimuli appeared as a hill-shaped mound rising from the horizontal plane of a receding surface. The concave stimuli appeared as a complementary depression recessed below the surface. The main source of shading information on these surfaces at any time was from one of four spotlights located at the diagonal corners of scene (see Fig. 1). Between trials, the direction of this light source was varied among these positions to create a wide variety of shading patterns on these surfaces. This was done to encourage the birds to perceive the generalized shapes of these surfaces rather than memorize specific lighting patterns. A second source of illumination was an elevated central overhead light that provided weak, even and diffuse illumination across the scene in order to make all aspects of the surface visible. To encourage further the generalized recognition of the surface shapes, two other irrelevant scene features were varied. Because birds experience scenes from different heights during flight, we manipulated the perspective of the surface shapes by rendering them from two camera




Fig. 1. Examples of the convex and concave surfaces tested in Experiment 1. The columns show the four different lighting directions used to produce differential surface shading across the surfaces. These examples show the 25° camera perspective and blue hue. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

viewpoints positioned at either 25° or 35° above the surface. Further, the surfaces were colored to have an overall blue or orange hue to create variety and to emphasize that the pattern of shading was the invariant feature to process.

Finally, we also decided to include both surface shading and shadow information in the displays in order to best simulate what the birds naturally experience. A feature of the software used to create these stimuli permitted us to independently include or remove the shadow cast by the surface shape. We subsequently discovered, however, that this clever feature permitted an error that was present for several of the initial experiments. Specifically, for those scenes illuminated from the front two lighting sources, the software switch responsible for producing shadows had been left off when those stimuli were rendered. As can be visually confirmed in Fig. 1, this oversight was easy to miss because these absent shadows are not readily visible as they are occluded by the shape of the surfaces (see 45° and 315° examples in Fig. 1). To our relief, it appears the pigeons never noticed either. After looking at the entirety of our results with this issue in mind, it appears it had little impact on their behavior. Nevertheless, where appropriate, we point to any issues that might have stemmed from this unintended error. It was immediately corrected upon its discovery in Experiment 4.

Together these manipulated features created 32 different surfaces that were tested (two shapes × four light sources × two perspectives × two hues). Using these differentially illuminated surfaces, the pigeons were tested in a go/no-go procedure in which they had to discriminate the convex or concave direction of the surface's shape. The first phase of Experiment 1 examined the acquisition of this discrimination, followed by a second phase that examined a block of sessions to evaluate post-acquisition performance.

2.1. Methods

2.1.1. Animals

Six male pigeons were tested. All had prior experience with different discrimination tasks involving rotating moving objects (Cook, Beale, & Koban, 2011; Koban & Cook, 2009). They were maintained at 80–85% of their free-feeding weight and individually caged in a colony room (12 h LD cycle) with free access to water and grit. All conditions and procedures followed guidelines and standards set by Tufts University IACUC.

2.1.2. Apparatus

Testing was conducted in a computer-controlled flat-black Plexiglas chamber. Stimuli were presented on an LCD color monitor (NEC AccuSync LCD51VM-BK; Wooddale, IL; not gamma corrected) visible through a 29 × 21.5 cm viewing window in the middle of the front panel of the chamber. Pecks to the monitor (1024 × 768 pixels) were detected by an infrared touch screen (Carroll Touch Systems, distributed by Tyco Electronics). The viewing window's bottom edge was 14.5 cm above the chamber floor. The monitor was recessed back 7.5 cm from the front panel. Mixed grain was delivered through a centrally placed food hopper (Coulbourn Instruments, Whitehall, PA). A 28-volt houselight was centrally located in the chamber ceiling and was always illuminated, except for timeouts.

2.1.3. Stimulus conditions

All stimuli were 480 × 360 pixels in size tested in a landscape orientation and generated by 3DS Max (version 7; Autodesk, San Rafael, CA, USA). The concave and convex surface shapes were created using the ripple modifier on a standard box shape to produce a smoothly curved concave or convex surface. To provide a framework for measurement within the scenes, the peak height to surface distance of the convexity was set to 70.5 units in the software and which we will refer to as a standard value of 1 for describing the remaining metrics of the scenes.

The shading in the scenes were generated using the built-in implementation of the Blinn modification of the Phong reflection equation (Blinn, 1977), with the shadows generated by a projective shadow mapping algorithm (Williams, 1978). The surface material for these surfaces were created from two different colors consisting of blue (base RGB values = 88, 199, 225) or orange (base RGB values = 227, 152, 152) hues, modified to include a modest amount of specularity (Phong specular factor; software's materials specularity parameter = 32) and glossiness (Phong shininess factor; software's materials glossiness parameter = 20). The illumination of each scene originated from two lighting sources. The first and most important were from a set of four omnidirectional lights (software's lighting intensity parameter = 1) located in the four diagonal corners of the scene relative to the observer's or camera's location: right-front (45° to the right of camera), right-rear (135°), left-rear, (225°) and left-front (315°). It was these lighting sources that created the differences in

surface shading, specularity, and casted shadows among the scenes. Their heights were 7 standard units above the surface and each positioned 7 standard units from vertical midline of the surface. Only one of these light sources was illuminated at a time. The second source of illumination was a dim omnidirectional light (intensity = 0.1) located 17 standard units above the surface and provided an even diffuse illumination of the entire scene. Finally, as noted in the introduction, each light had a software switch for controlling whether or not a shadow was cast by this lighting source. Because its results were not readily visible, this switch was unintentionally left off for the two front lighting directions. This switch was set for the rear two lighting directions and they produced normal and readily visible shadows.

All surfaces were rendered from two camera perspectives. These perspectives were 4.7 standard units from the scene's central point at either a 25° or 35° angle above the surface. Each rendering computed the perspective projection from the position of the camera with a field of view of 39.5 degrees aimed at the center of the scene. All together these four factors (two surface shapes, two colors, four lighting direction and two camera angles) resulted in 32 different stimuli that were tested in Experiment 1.

2.1.4. Procedure

2.1.4.1. Preliminary training. Each session consisted of 96 trials. Following a peck to a 2.5 cm white warning signal, a surface was presented for 20 s in the center of the screen. These 96 trials consisted of 48 convex and 48 concave presentations with each of the 32 surface stimuli tested three times. The order of presentation was randomly determined for each session. Pecks to all stimuli were non-differentially reinforced on a variable-interval schedule (VI-10) with access to food for 2.9 s. Because the pigeons were familiar with pecking the screen, little preliminary training was needed. Once the pigeons' peck rates stabilized to these stimuli over two or three sessions, discrimination training commenced.

2.1.4.2. Discrimination training. Except for the introduction of differential reinforcement depending upon the surface's shape, discrimination testing continued in the same way as preliminary training. For the three pigeons in the concave+ condition, pecks to the concave surface resulted in food reward (VI-10 schedule), while pecks to the convex surface led to a dark timeout proportional to the number of pecks (1-s per peck). At the end of the S+ presentations an additional reward was automatically delivered to the pigeons. For the three pigeons in the convex+ condition, these reward contingencies were reversed. To allow for the uncontaminated measurement of peck rate without the disruption of food presentations, 16 non-reinforced *probe* trials of the positive condition were conducted each session. These sixteen randomly-placed probe trials tested each combination of lighting direction, camera perspective, and color. All dependent measures were calculated from these probe trials. Discrimination training was conducted until the pigeons' performance stabilized with three or more sessions with a discrimination ratio (DR = S+ probe peck rate / [S+ probe peck rate + S- peck rate]) of .7 or

greater. At this point, 16 additional sessions of testing using the same session organization were conducted to examine baseline discrimination.

2.2. Results

2.2.1. Acquisition

Overall, the discrimination was easily learned by the pigeons. Fig. 2 shows the acquisition of each bird across sessions as measured by DR. Four pigeons were showing unambiguous evidence of discriminative behavior by the end of the third (<288 trials) or fifth (<480 trials) session. Three pigeons were even showing the beginnings of discrimination by the end of the first session. One pigeon was considerably slower than the others, needing approximately 20 sessions to reach a level of performance that the others exhibited in 10 or fewer sessions. As determined from the last session of acquisition, all six pigeons (mean DR = .82,) were significantly above chance performance, $t(5) = 7.0$ (DR = .5; an alpha level of .05 or less was used to judge this and all subsequent statistical tests).

The convex+ discrimination was possibly learned slightly faster than the concave+ discrimination. As measured by mean sessions to criterion (three sessions of DR > .7), the three pigeons in convex+ group numerically learned their task (mean = 7.7 sessions) more quickly than the three pigeons in the concave+ group (mean = 15 sessions). A statistical comparison is difficult, however, because of the one slow bird, but the pigeons in the concave+ group were the third, fourth and slowest birds during learning.

2.2.2. Post-acquisition behavior

Following the completion of acquisition, the pigeons were discriminating the surface shapes very well. Looking at total pecks over the presentation interval, all six birds were pecking at higher rates to their positive stimuli (S+ = 29.5) in comparison to the negative ones (S- = 10.3). This was true for each group (concave+ S+ = 32.3, S- = 12.6, convex+ S+ = 26.6, S- = 8.1). Single mean *t*-tests confirmed that each of the six pigeons were significantly discriminating between their convex and concave surfaces, all $t_s(15) > 2.1$. We next examined the effects of the different surface colors, lighting directions, and camera perspectives on discrimination. Overall, these variations among the stimuli had very little impact on discrimination. A repeated measures ANOVA (Positive/Negative Surface Direction [adjusted for group] × Surface Coloration × Perspective × Lighting Direction) using total pecks was conducted. Not surprisingly, this analysis revealed a significant effect of positive vs. negative surface direction, $F(1,5) = 223.3$, partial $\eta^2 = .98$. Surface direction did not significantly interact with either lighting direction (45°-S+ = 29.8, S- = 11.0; 135°-S+ = 29.0, S- = 9.7; 225°-S+ = 29.4, S- = 9.4; 315°-S+ = 29.6, S- = 11.1; $F(3,15) = 1.7$) or surface coloration (orange-S+ = 29.4, S- = 9.9; blue-S+ = 29.6, S- = 10.8, $F(1,5) = 1.2$) as peck rates over these factors were equivalent. There was a significant interaction between surface direction and perspective, $F(1,5) = 17.3$, partial $\eta^2 = .78$, as the 25° condition (S+ = 29.7, S- = 9.9) supported a consistent, though small, advantage over the 35° condition (S+ = 29.3,

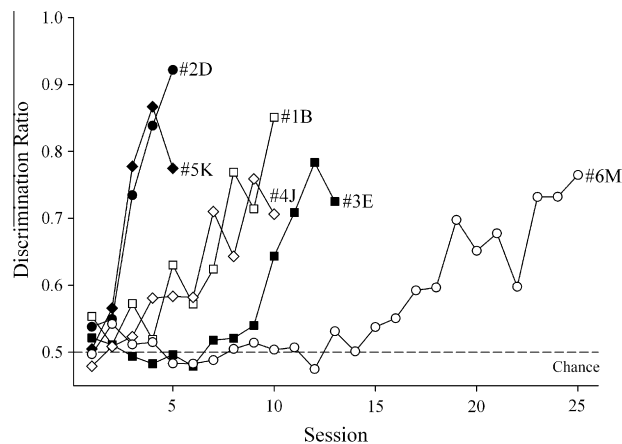


Fig. 2. Acquisition of the discrimination of surface direction as a function of discrimination ratio for the six pigeons tested in Experiment 1 (filled symbols = convex+ group, open symbols = concave+ group). The dotted reference line shows chance performance.

$S- = 10.7$). This analysis also revealed a significant main effect of lighting direction, $F(3, 15) = 5.4$, partial $\eta^2 = .25$, as the front lit surfaces (mean = 15.7) generated slightly more overall pecking than the back lit conditions (mean = 14.5).

Our next analyses were directed at determining how quickly the pigeons detected the direction of the surface's shape. Here we divided the number of pecks within a presentation as a function of when they occurred. Shown in Fig. 3 are mean number of pecks for the $S+$ (adjusted for direction across groups) and $S-$ stimuli as a function of presentation time in 250-ms bins. The pigeons very quickly showed differences in their peck rate depending upon surface direction. Using a series of paired t -tests comparing each point along the two functions, the pigeons begin to significantly differentiate among the $S+$ and $S-$ stimuli at 750 ms after presentation, $t(5) = 3.6$. The rapidity of this discrimination within a trial is consistent with the ease of learning the discrimination found more generally across sessions.

We further examined performance across these temporal bins in a repeated measures ANOVA (Surface Direc-

tion \times Surface Coloration \times Lighting Direction \times Perspective \times Time [but in 500 ms bins to accommodate an SPSS constraint]). Besides the same effects described previously, this analysis confirmed the interaction between surface direction and time, $F(39, 195) = 16.0$, partial $\eta^2 = .76$. In addition to the expected difference early within a presentation as peck rates diverged, we also observed a difference at the end of the presentation as the relative size of the discrimination declined. In particular, three of the six birds started to show an increase in pecking at the negative displays near the end of the interval. One speculation is that it may reflect some form of temporal generalization across all conditions because of the "free" time-based reward at the end of the $S+$ trials. Regardless, to ensure a fair and accurate measure of the discriminative behavior of all six pigeons, we used the peck rates over the first 5 s of each presentation to best measure discrimination in future analyses.

We next evaluated the potential role of overall illumination, the amount of black, or the amount of white in the stimuli and their relations to discrimination. These

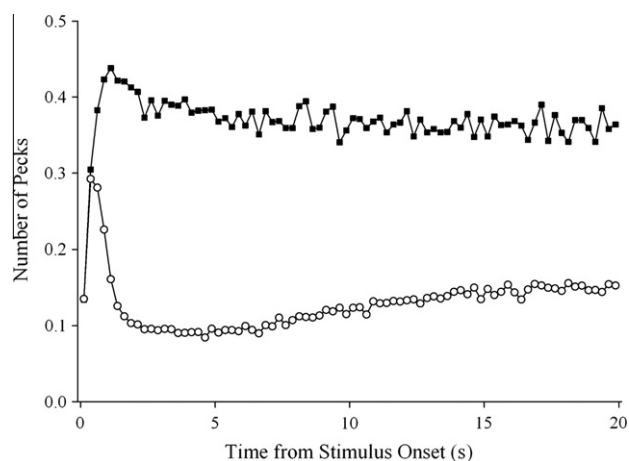


Fig. 3. Mean number of pecks across presentation time in 250-ms time bins for $S+$ and $S-$ conditions during post-acquisition testing in Experiment 1.

components were measured by converting all 32 images to gray-scale values. Overall illumination was calculated from the mean pixel value of the entire image. The amount of black and white was computed by normalizing the pixel values to the maximum and minimum values shown,

dividing these gray-scale values into eight equal-sized bands, and then counting the number of pixels in the lowest and highest bands. Fig. 4 shows for each group the mean number of pecks over the first 5 s as a function of these descriptive measures. Overall, these simple image

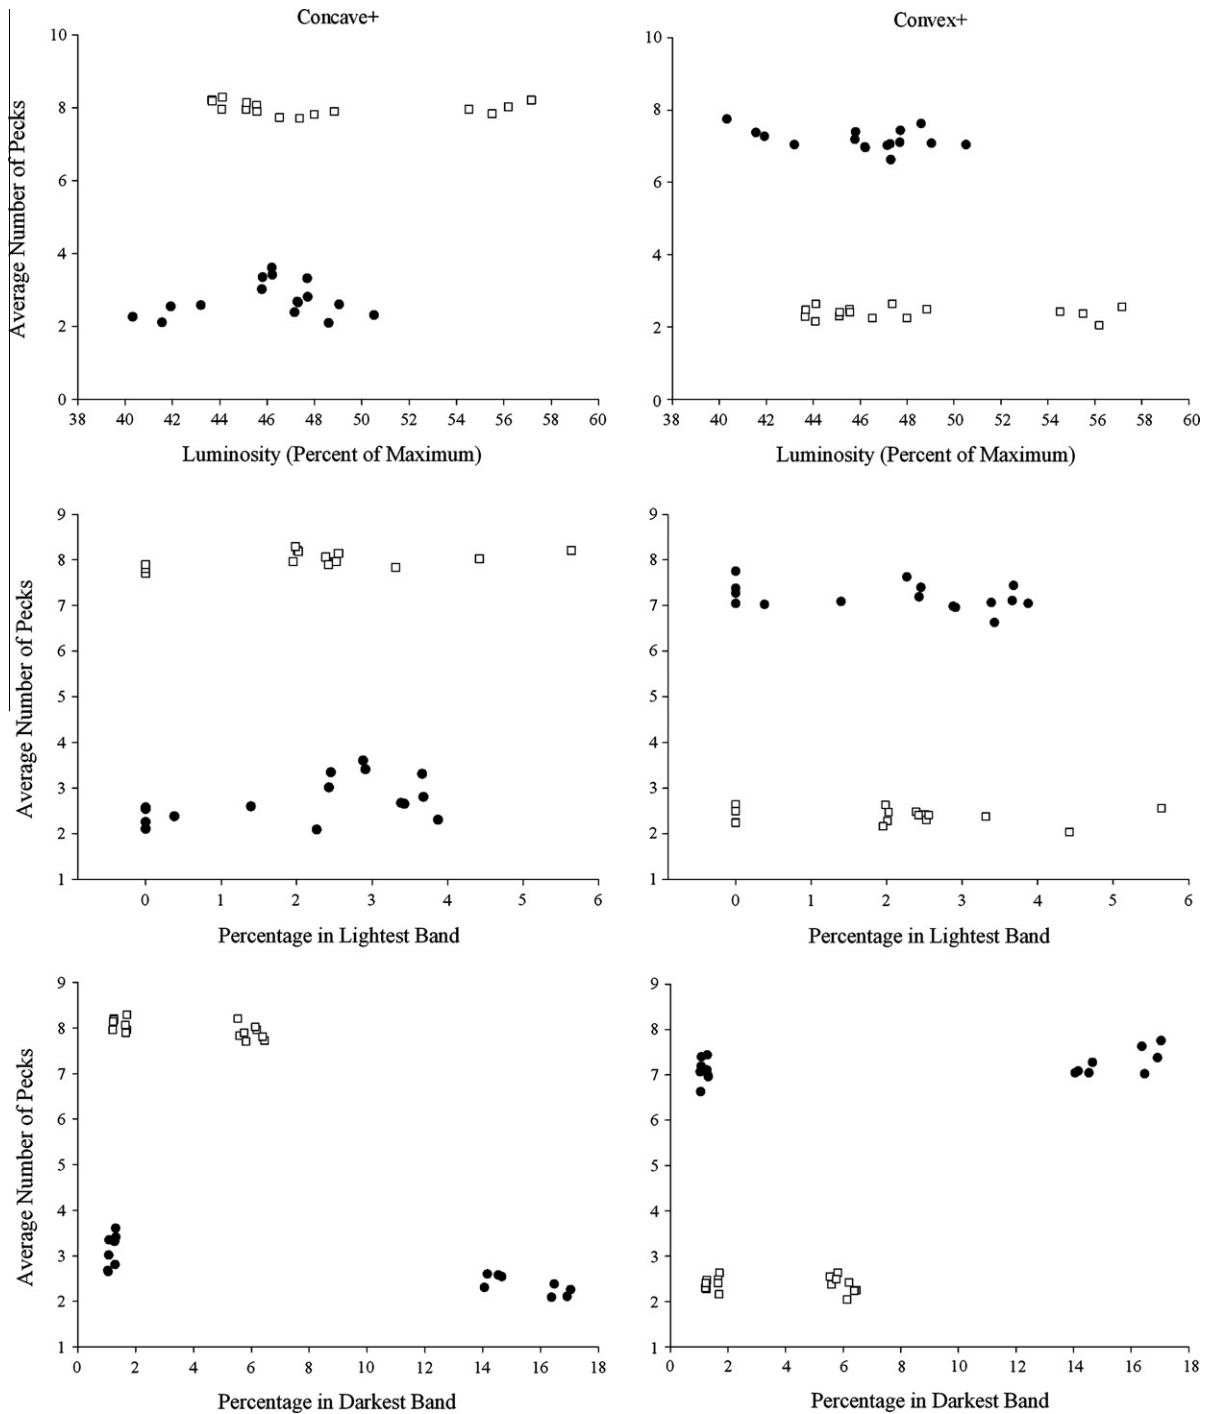


Fig. 4. Mean number of pecks in the first 5 s as a function of stimulus luminosity, lightest and darkest gray-scale bands during post-acquisition testing in Experiment 1. Columns show performance for each group. Filled black symbols show the values for 16 convex surfaces and open square symbols show the values for 16 concave surfaces for each group.

measures failed to capture the robust discriminative behavior observed in either group indicating that these features were not responsible for the discrimination.

2.3. Discussion

The pigeons quickly learned to discriminate the direction of the convex and concave surface shapes, with the majority of birds demonstrating evidence of learning within just a few sessions. Consistent with the ease of learning, discriminative behavior also emerged quickly upon stimulus presentation, with significant differences in peck rate occurring by 750 ms after stimulus onset. Overall, the convex+ group possibly had a slight learning advantage over the concave+ group, albeit this is complicated by the one much slower learning bird in the latter group. Irrelevant variation in lighting direction, surface coloration, and camera perspective had little impact on discrimination. The mistaken absence of projected shadows in the front-lit scenes apparently made no difference, as the pigeons discriminated the surfaces from the different lighting directions about equally well. Overall, the results suggest that the differential lighting on the surfaces provided the pigeons a set of highly salient visual cues that could be used to differentiate the convex or concave direction of the surface shape.

These results are thus consistent with the hypothesis that differential lighting on the shapes of these surfaces plays an important role in avian visual processing. This differential lighting resulted in several visual features that could have been useful. One would be the graded shading on the curvature of the surfaces. As noted, such shape-from-shading cues are important to object recognition and depth processing in humans. Another visual feature created by the differential lighting was the contrast contours present in some of the images. These contours were produced by the contrast in light reflecting from the front of the convex shape and its interposition against the contrasting background illumination of the receding surface. A weaker contrast contour exists for similar reasons along the lower edge of the concave shapes. These contours could have contributed to the ease of the discrimination by helping to promote the appearance of depth within the scenes.

A key question concerns what perceptual and representational factors produced this rapid and robust discriminative behavior based on illumination. These shaded stimuli generate a strong impression of depth-based surfaces that allow humans to easily see the 3D convexity and concavity of the shapes portrayed in them. Is this the case for the pigeons, too? When we examined several simpler image-based factors in the stimuli, they did not account for peck rate differences across the differentially illuminated concave and convex surfaces. This suggests that it was indeed the relative pattern of light and dark within the images that was most important.

Given this, there are two possibilities that can be entertained. The first alternative is that the pigeons do not see the intended shape and depth in these stimuli, but instead use the 2D pattern of luminance in the stimuli as the basis for their discrimination. This hypothesis suggests the pigeons were not seeing the convexity and concavity of the shapes (since these are depth-based features), but learned and

reacted to the image-based patterns of light and dark textures experienced during training as the basis for their discrimination. This alternative implies that the front-lit and back-lit scenes are likely represented as different luminance patterns since these stimuli share little in common without seeing the underlying 3D shape. Given what is established about their large memory capacity (Cook et al., 2005; Fagot & Cook, 2006), having such multiple and independent representations would not be particularly challenging.

The second alternative is that the pigeons are experiencing these stimuli as intended. In this case they are visually processing the patterns of light and dark shading to generate 3D shape-based perceptions of scenes as portrayed in depth. Thus, the pigeons are able to extract the convexity and concavity of the shapes regardless of illumination direction and discriminate among the stimuli based on this shape-based feature. The ease and invariance of the discrimination across the irrelevant features of surface coloration, camera perspective, and lighting direction seems most consistent with this hypothesis. It might also account for the slight learning advantage of the convex+ group. This group would have a surface “feature” in the form of a projected shape to peck at that would be missing for the concave+ group (Jenkins & Sainsbury, 1970; Wasserman & Anderson, 1974). Nevertheless, results of Experiment 1 cannot strongly differentiate among these competing alternatives of the pigeons’ perceptual experience. These alternatives do, however, provide the motivation for the remaining experiments.

3. Experiment 2

Experiment 2 consisted of two tests to better understand how the surface shapes were perceived. The first examined how the pigeons transferred and reacted to changes in the camera’s viewpoint relative to the surfaces. This served to examine the generality of the discrimination learned in Experiment 1. Such transfer would indicate that it was the relative, rather than absolute, pattern of lighting on the surfaces that was important. If the pigeons had learned to discriminate the shape of the surfaces, then they should show transfer over a reasonable range of camera perspectives. Only as the camera’s viewpoint becomes increasingly perpendicular to the surface and the 3D shape of the surfaces gradually disappear should discrimination deteriorate. On the other hand, if the pigeons were just experiencing 2D luminance patterns then any change in viewpoint should affect performance, since movement in either direction causes substantial variation in the pattern of shading from that experienced during training.

The second test investigated how the vertical extent of the convexity and concavity of the surface shapes influenced the discrimination by varying the height and depth of the surface shapes. If pigeons were attending to the shape of the surface, then changing the height of this shape should influence the discrimination. We expected to see a stronger effect of these manipulations with the convex surfaces, since such vertical changes would be most readily visible. Although the same changes were rendered with the concave surface shapes, their effects would be more difficult to detect as they were visually occluded and harder to see.

3.1. Methods

3.1.1. Animals and apparatus

The same six pigeons and apparatus were used as the previous experiment. After completion of the perspective test, one pigeon did not further participate for reasons unrelated to the experiment.

3.1.2. Procedure

3.1.2.1. Variation in viewpoint. The baseline stimuli for this test were the same as in Experiment 1. They involved 96 baseline trials (48 S+/48 S-) testing equal numbers of the two camera perspectives, four lighting directions, and two surface colorations. One third of the positive trials were probe trials.

The test stimuli were generated by varying the camera's angular perspective relative to the surface in 10° increments from 5° to 85°. These stimuli were generated for each combination of surface color (blue and orange) and two lighting directions (45° and 135°). This resulted in 56 test stimuli (seven novel test angles × two surface shapes × two colors × two lighting directions). A subset of these stimuli was randomly inserted in the daily baseline sessions as sets of 14 non-reinforced probe trials. Each resulting 110-trial session tested all novel angles with each surface at a different combination of color and lighting direction. Eight test sessions were conducted.

3.1.2.2. Variation in surface height and depth. The baseline stimuli for each session were the same for this test, but organized slightly differently within a session. Only 84 baseline trials (42 S+/42 S-) were tested each session with equal numbers of the two camera angles, but with lighting direction and surface color now randomized across trials within a session. Ten percent of the positive trials were conducted as probe trials.

The test stimuli were generated by varying the height and depth of the surface shape as a percentage of the standard unit (150%, 50%, 25% and 0%). These stimuli were generated for only the blue surface color (done to reduce the amount of rendering time and probe testing, since this feature was having no effect), all four lighting directions, and both baseline camera perspectives (25° and 35°). This resulted in 64 test stimuli (four heights/depths × two surface directions × four lighting directions × two camera angles). A subset of these stimuli was randomly inserted in the daily baseline sessions as a set of 16 non-reinforced probe trials. Each resulting 100-trial session tested all novel heights/depths at each angle in combination with a single lighting direction (the latter in a series of 4-session blocks). Sixteen test sessions were conducted.

3.2. Results

3.2.1. Variation in viewpoint

The two panels of Fig. 5 shows the results of varying the camera's perspective for the convex+ and concave+ groups based on the mean number of pecks in the first 5 s of presentation. All the birds showed solid discrimination transfer across an approximately 60° range of camera perspectives, with a clear separation between S+ and S- peck

rates from 5° to 65°. As the camera's viewpoint became increasing perpendicular to the surfaces, discrimination systematically declined as the 3D shape of the surfaces became increasingly more difficult to discern. Within each panel, the results are also separated by whether the surfaces were illuminated from the front or back lighting direction. Overall, the pattern of transfer for each group was highly similar regardless of illumination source's location.

We then conducted a repeated measures ANOVA (Surface Direction [adjusted for group] × Camera Perspective × Lighting Direction × Surface Color) using the average number of pecks over the first 5 s. The two training perspectives were excluded from this analysis. It confirmed the interaction clearly present in Fig. 5 between surface direction and camera perspective, $F(6,30) = 14.0$, partial $\eta^2 = .74$. Neither surface color nor lighting direction had significant interactions with surface direction. As the latter factors were not critical, we collapsed peck rates across them and examined performance at each perspective with a series of paired *t*-tests, critical value $t(5) = 2.57$. All six pigeons significantly discriminated surface direction at each perspective, except in the 85° condition (S+ = 4.3, S- = 4.1).

3.2.2. Variation in surface height and depth

The two panels of Fig. 6 shows the results of varying the height and depth of the surface shape for the convex+ and concave+ groups based on the mean number of pecks in the first 5 s. The pigeons showed excellent transfer to conditions having either a 50% increase or decrease from the height of the standard training surface. The 25% condition began to show the impact of the reduced height, with the flat surface not supporting the discrimination (see Fig. 7).

These results were confirmed by a repeated measures ANOVA (Surface direction × Surface Height × Camera perspective × Lighting Direction) based on the first 5 s of pecking in each condition. As to be expected from Fig. 6, surface direction significantly interacted with surface height, $F(3,12) = 42.5$, partial $\eta^2 = .91$. This interaction was further modified by its significant interaction with lighting direction, $F(9,36) = 2.7$, partial $\eta^2 = .41$. The source of the latter triple interaction was due to a differential effect of lighting at the 25% height condition. In this case, both groups of pigeons discriminated the front-lit surface better than the back-lit one. Based on paired *t*-tests, the front-lit 25% condition supported significant discrimination, $t(4) = 3.2$, albeit at a level reduced compared to taller surfaces (S+ = 4.2, S- = 2.0). Although about the same sized difference (S+ = 3.8, S- = 1.9) the back-lit condition did not support significant discrimination, $t(4) = 2.0$. Identical comparisons for 50% and 150% conditions revealed highly significant discrimination in both groups and no differential effect of lighting direction.

3.3. Discussion

Experiment 2 showed the pigeons' capacity to generalize their discrimination to a wide range of novel perspectives on the surface shapes, but that declined as the perspective became increasingly perpendicular to the

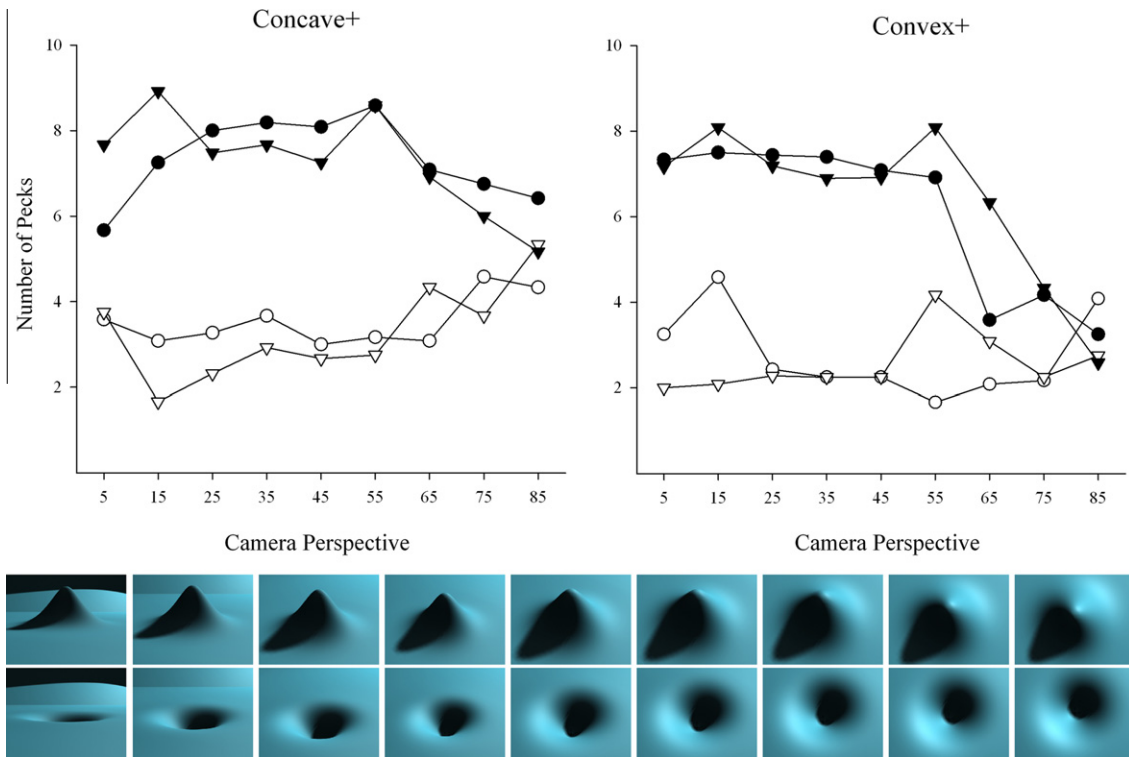


Fig. 5. Mean number of pecks in the first 5 s as a function of camera perspective for each group in Experiment 2 (filled symbols = S+, open symbols = S-). The shape of the symbols mark which lighting direction was tested (circle = 45°, triangle = 135°). Examples of the stimuli tested are included below the figure.

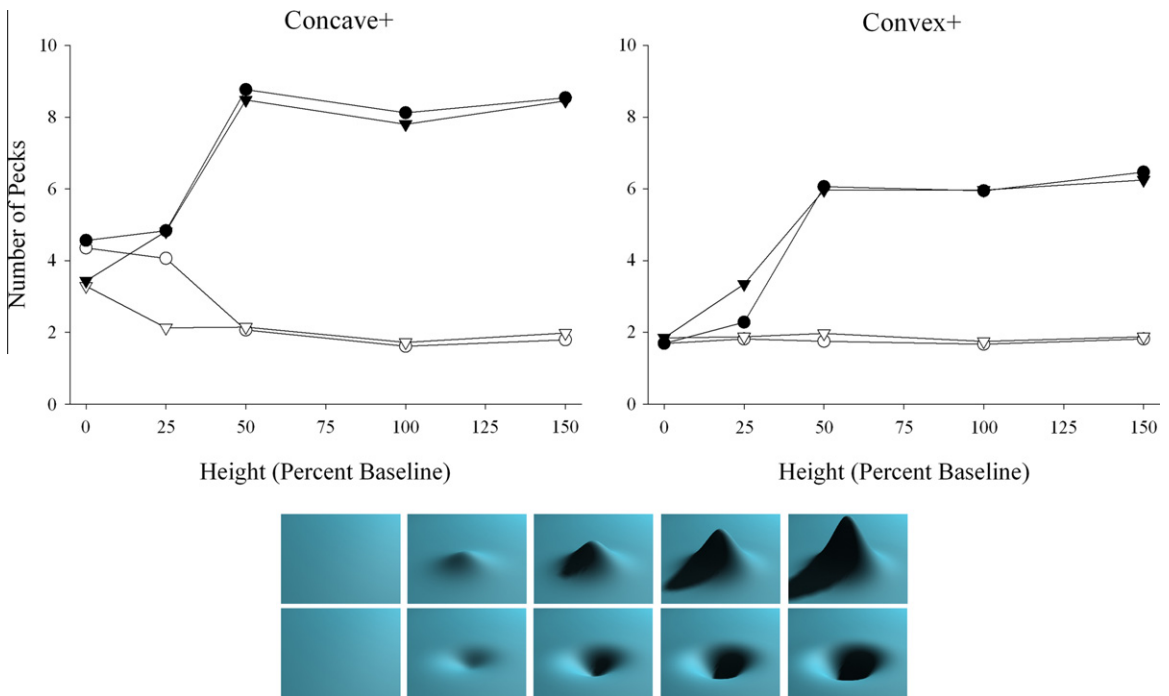


Fig. 6. Mean number of pecks in the first 5 s as a function of surface height/depth in Experiment 2 (filled symbols = S+, open squares = S-). The shape of the symbols mark the lighting directions tested (circle = 315° and 45°, triangle = 225° and 135°). Right and left panels show performance for each group. Examples of the stimuli tested are included below the figure.

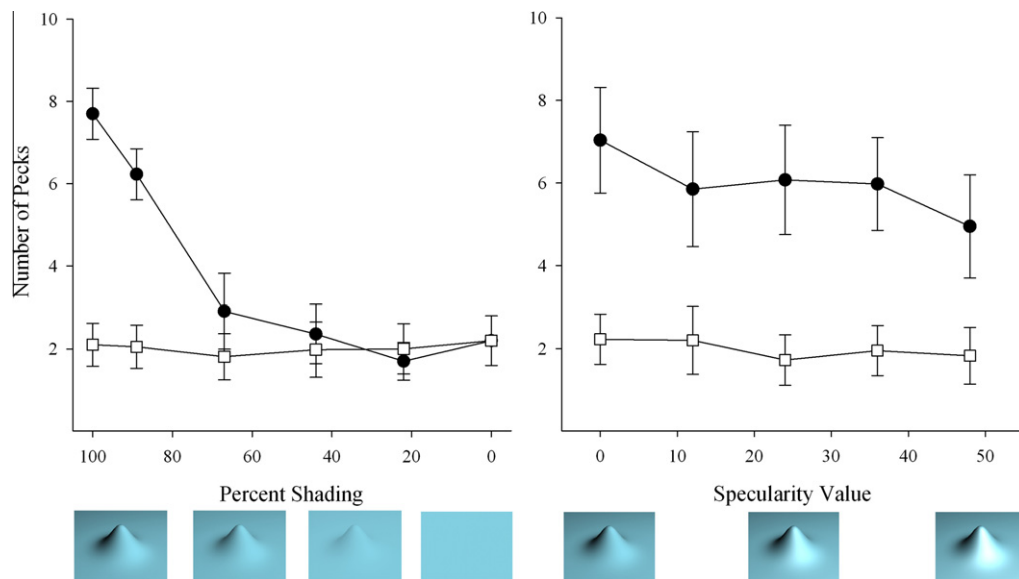


Fig. 7. Mean number of pecks in the first 5 s across all birds for the two tests in Experiment 3 (filled circles = S+, open squares = S-). The left panel shows performance as a function of relative contrast. The right panel shows performance as a function of specularity value used to render the different surfaces. Selected examples of the stimuli tested are included below the figure.

surface. The pigeons also transferred to novel variations of the original height and depth of the surface shapes that was limited by the vertical extent of this change. The overall pattern of these results indicates that the pigeons were likely discriminating the implied concavity or convexity of the surfaces as provided by the lighting cues within the scenes. This is suggested by the fact that the decline in the pigeons' performance was not affected by just any variation in the patterns of shading, as predicted by the generalization of an image-based 2D account, but instead was tied to the specific elimination of shaped-based information as derived from either reduced height or vertical perspective. Increases in height or lowered perspectives that retained shape-based information did not bother the pigeons, despite the fact that these changes all produced changes in the patterns of shading.

Overall, the pigeons showed an ability to discriminate these surface shapes over a range of 20–25° beyond the values experienced in training. Although difficult to directly compare, this value is generally in line with previous experiments looking at generalization to novel views along different axes as tested with asymmetrical objects (Peissig, Wasserman, Young, & Biederman, 2002; Spetch & Friedman, 2003). As the camera's viewpoint became increasingly overhead, the pigeons' performance gradually declined. At least two reasons for this latter decline are possible. One is that the three-dimensional features for determining surface height become difficult to see as the shape of surfaces becomes increasingly flat and two-dimensional. Birds in flight might experience similar difficulties as they move over natural objects and experience less canonical views of them. A second possibility is that there is some inherent limitation in the degree to which rotated objects are recognized by pigeons because they do not match the shape representation stored from previous encounters. The present

data cannot decide between these accounts. Future research using asymmetrical surface shapes and varying the number and types of views experienced during training should help to resolve this question and whether our scene-based results are similar or different from those previous tests using isolated objects.

The same issue is present regarding the effects of surface height. Again this could be due to a perceptual limitation, as the relative convexity and concavity of the surfaces shrinks and disappears and the birds simply cannot see the shapes required for the discrimination. Likewise, however, these changes are further away from the training value and may represent a limitation on the degree of shape recognition the birds can exhibit from previously learned experiences.

4. Experiment 3

The goal of Experiment 3 was to understand how the relative salience of light and dark regions in the scenes influenced the discrimination. The experiment examined two changes in the properties of the surfaces that altered the character of the reflected illumination. In the first test, we reduced the contrast of the light and dark regions comprising the surfaces. If the relative shading in the scene was critical to the discrimination, then gradually eliminating this contrast should reduce and eliminate the discrimination.

In the second test, we manipulated the specularity of the light reflected from the surfaces. The reflectance of different materials is a combination of both diffuse and specular components. Highly reflective materials, for example, have strong specular highlights consisting of the bright whitish points at the convexities and contours of their shape. This type of reflectance provides a valuable cue to the object's 3D shape and the location of the light source.

It is important in the human perception of scenes and is a critical feature in generating realistic-looking computer graphics of objects and scenes. If the birds could continue to perform the discrimination despite substantial alterations in the specular pattern of lighting and shading, it would support the idea that they were using the lighting to primarily extract the 3D structure of these surface shapes rather than relying on memorized 2D patterns of lighting and shading.

Prior to this experiment, several changes were made to the pigeons' daily training. These consisted of testing the concave and convex surfaces from the 25°, 35°, 45°, 65°, and 85° camera perspectives used in Experiment 2. The inclusion of these different perspectives was to continue to encourage the pigeons to recognize and represent these surface shapes independently of their specific pattern of illumination. Because surface coloration had no detectable effects in the first two experiments, we eliminated the orange surfaces from further testing. Finally, the four different lighting directions were now selected at random across trials rather than being tested equally often within a session.

4.1. Methods

4.1.1. Animals and apparatus

The five pigeons and apparatus were the same as completed Experiment 2.

4.1.2. Procedure

4.1.2.1. Contrast variation. This test varied the relative amount of the light and dark in the scenes. This contrast was changed, while maintaining overall luminosity, by reciprocally altering the "brightness" values of the ambient color and self-illumination software switches controlling image rendering. This gradually reduced the amount of apparent shading on the surfaces as a percentage of original shading (88%, 66%, 44%, and 22%). Because lighting direction had not shown any differential effects, test stimuli were only rendered and tested from the 45° lighting direction. Four camera perspectives were tested: 35°, 45°, 65°, and 85°. This resulted in 32 test stimuli (four levels of contrast \times two surface shapes \times four camera angles). A flat surface similar to that used in Experiment 2 tested at the 35° viewpoint was also included in testing.

The test stimuli were randomly inserted as non-reinforced probe trials into daily baseline sessions. The baseline sessions consisted of 100 trials (50 S+/50 S– with / 10% S+ probe trials), testing equal numbers of blue-hued surface shapes from among five camera perspectives (25°, 35°, 45°, 65°, and 85°) with the same four lighting directions now randomized across trials. Each 125-trial session tested two of the four contrast values at each camera angle, the flat stimulus, and a set of normally shaded 100% values at each camera angle. A total of six of the resulting test sessions were conducted with each contrast level tested equally often.

4.1.2.2. Specularity variation. The second test varied the amount of reflective specularity in the material used to render the surface shapes. This was manipulated by chang-

ing the value of the specularity parameter in the rendering program (48, 36, 23, 12, 0; baseline software value = 32). As with the contrast test, test stimuli were only tested with lighting from the right-front direction. Four camera angles were again tested: 35°, 45°, 65°, and 85°. This resulted in 40 test stimuli (five levels of specularity \times two surface shapes \times four camera angles).

These test stimuli were randomly inserted as a set of non-reinforced probe trials into daily baseline sessions. The baseline trials were organized in the same way as described for the contrast test. Each 124-trial session tested three values of altered specularity stimuli at each camera angle. A total of six test sessions were conducted with each specularity level tested equally often, except for the zero value which was tested each session.

4.2. Results

4.2.1. Contrast variation

As predicted, the presence of shading controlled the discrimination; performance declined as this information faded away. The left panel of Fig. 6 contains the mean number of pecks during the first 5 s of stimulus presentation for the different contrast conditions and shows the reduction in discrimination with contrast removal. Neither variation in the camera perspective nor convex+ or concave+ assignment had an effect on this pattern.

A repeated measures ANOVA (Surface Direction \times Contrast \times Perspective) on the number of pecks in the first 5 s of presentation confirmed this outcome. The zero contrast condition was excluded from this analysis because it was only tested at one camera perspective. As expected from the functions displayed in Fig. 6, this analysis found a significant interaction between surface direction and contrast, $F(4, 16) = 39.4$, partial $\eta^2 = .91$. It also revealed a significant main effect of contrast, $F(4, 16) = 37.4$, partial $\eta^2 = .90$, consistent with the overall decline in pecking observed with the reduction in contrast. There were no significant main effects or interactions with camera perspective among these different factors.

4.2.2. Specularity variation

Changing the specularity of the material in the surface scenes had little effect on discrimination performance. The pigeons continued to discriminate among the convex and concave surfaces despite dramatic changes in the shaded appearance of the stimuli caused by this manipulation. Shown in the right panel of Fig. 6 is the mean number of pecks during the first 5 s of stimulus presentation for the different specularity conditions. The greater levels of specularity change did appear to reduce the performance slightly, but in general the pigeons continued to perform well. Again, variation in the camera perspective had little effect on performance.

These general patterns were confirmed by a repeated measures ANOVA (Surface Direction \times Specularity \times Perspective) on the number of pecks in the first 5 s of presentation. This analysis found a significant main effect of surface direction, $F(1, 4) = 119.5$, partial $\eta^2 = .97$, that extended across all conditions. There was no significant interaction between surface direction and specularity,

$F(4, 16) = 1.8$. There was a significant main effect of specularity, $F(4, 16) = 4.6$, partial $\eta^2 = .53$, which resulted from the slightly reduced levels of pecking with increasing specularity independent of surface direction. There were no significant main effects or interactions with camera perspective for these different factors.

4.3. Discussion

Experiment 3 revealed how changing two important properties of illumination related to the surfaces affected discrimination. First, the systematic reduction in contrast correspondingly reduced the pigeons' capacity to perform the discrimination. This confirms that it was the contrasting pattern of light and dark in the stimuli that is critical. When the specularity of the surface's material was varied, resulting in considerable changes in the patterns of reflected light from the surfaces, the discrimination remained strong. These experiments further confirm that across substantial transformations in the structure and organization of their relative shading, the specific pattern of light and dark is not particularly critical as long as shape-based information is present. When combined with the asymmetrical transfer across perspectives in Experiment 2, it continues to suggest that the pigeons were capable of recognizing the surface shapes as long as they conformed to the structure of a familiar 3D shape.

5. Experiment 4

Is the familiarity of this shape critical? Experiment 4 examined how the pigeons transferred this shape-from-shading discrimination to novel surface shapes. If the pigeons were indeed seeing and responding to the concave and convex nature of the original surface's 3D shape, then they should exhibit some degree of discrimination transfer to these same features when presented as novel surface shapes. Given the stark differences in the patterns of light reflected by each shape, such discrimination would further indicate that the image patterns are not the critical cue for the discrimination.

This experiment had three phases. In the first phase, five different novel surface shapes were tested to investigate the generality of the already learned discrimination (see examples in Fig. 8). These novel surfaces were constructed with a number of different shape features that differentiated them from the shape of the original surface, including their degree of symmetry, relative height, and structural complexity. The idea was to test a variety of shapes to see if their shapes and the associated patterns of shading affected transfer. We hypothesized that transfer would be best to surfaces that were more similar to the shape of the training stimuli, such as the symmetrical circular "sphere" surface, and poorest to surfaces that looked most different, such as the asymmetrical, shallower, and multi-peaked, "bubble" surface.

The second phase of the experiment involved the differential training of a subset of these stimuli as part of their baseline sessions. The objective was to expand the number and type of surfaces experienced by the pigeons in order to

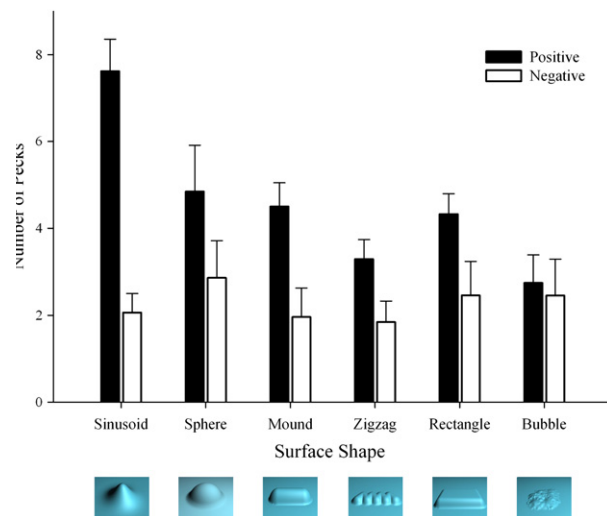


Fig. 8. Mean number of pecks in the first 5 s across all pigeons during transfer testing with the different novel surface shapes in Experiment 4 (filled bars = S+, open bars = S-). Examples of the stimuli tested are included below the figure.

reinforce the common feature of shape orientation as the most reliable feature for performing the discrimination. During this training only two lighting directions were used. This allowed us to test for shape "transfer" again by testing with the two remaining 90°-shifted lighting directions that had not been trained – and which completely inverted the pattern of lighting, but not the shape information, on these surface shapes.

The third phase trained and tested the birds with a newly rendered set of stimuli. It was at this point in time that we discovered the incorrect software setting for producing shadows from the front-lit light sources. While this had apparently not caused any consistent differences in their reactions to the prior stimuli, we nevertheless immediately stopped the experiment and rendered a new corrected set of stimuli. Thus, the third phase of the experiment consisted of the continuation with this amended stimulus set.

5.1. Methods

5.1.1. Animals and apparatus

The pigeons and apparatus were the same as in Experiment 3. During the break in testing, the touchscreen was replaced (EZ-170-WAVE-USB, EZscreen, Houston, TX).

5.1.2. Procedure

5.1.2.1. Novel surface shapes. Five novel convex and concave surface shapes were created (see Fig. 8, labeled sphere, mound, zigzag, rectangle and bubble; the original shape will be called sinusoid). The new shapes were created from a modified mesh box shape with the additional application of a non-uniform rational mesh smoothing operator to increase shape curvature. The concave conditions were again 180° inversions of the convex surfaces. The peak height/depth of each new shape in standard units was: sphere = .73, mound = .76, zigzag = .32, rectangle = .43,

bubble = .27. These surfaces were rendered using surface properties analogous to the original stimuli, although the spatial positioning of the sphere relative to the lighting sources resulted in it being rendered a little bit brighter than the other stimuli. Each new surface shape was rendered from the 45° lighting direction and the 35°, 45°, 65°, and 85° camera perspectives. This resulted in 40 test stimuli (five new surface shapes \times two surface directions \times four camera perspectives).

These test stimuli were randomly inserted as a set of eight non-reinforced probe trials into daily baseline sessions. The baseline sessions consisted of 100 trials (50 S+/50 S– with 10% S+ probe trials), testing equal numbers of the blue-hued sinusoidal shape from five perspectives (25°, 35°, 45°, 65°, and 85°) with the four lighting directions randomized across trials. Each 108-trial session tested a different novel shape at all camera angles. A total of 15 test sessions were conducted with each novel surface shape tested once in three randomly-ordered blocks.

5.1.2.2. Novel shape training and lighting transfer. After this, a subset of the novel shapes (sphere, mound, zigzag) were incorporated into the differentially reinforced baseline sessions. Each session lasted 82 trials and consisted of 58 trials (29 S+/29 S–) testing the original sinusoidal stimuli and 8 trials (4 S+/4 S–) each testing the three new surface shapes. The sinusoidal stimuli were tested from five camera perspectives (25°, 35°, 45°, 65° and 85°) and four lighting directions (45°, 135°, 225° and 315°) as randomly chosen each trial. For the three new shapes, these were restricted to a smaller set of randomly-chosen (perspectives – 35°, 45°, 65° and 85° and lighting directions – 45° and 225°). Twenty-five percent of the positive trials were randomly conducted as non-reinforced probe trials. This training lasted 42 sessions.

Probe transfer sessions were then conducted testing the three previously novel shapes as illuminated from the two 90°-shifted lighting directions (135° and 315°) not yet trained with these shapes. Each 88-trial session tested six additional randomly-inserted, non-reinforced probe trials testing each previously novel shape from one of the novel lighting directions. Each session tested one of the four different camera perspectives selected at random. The first four sessions tested the 135° lighting direction and the next four tested the 315° direction.

5.1.2.3. Amended set training. Upon discovering the error with the front-lit shadows, an entirely new set of stimuli were created. Besides including shadows for the front-lit scenes, we employed a more precise ray tracing algorithm for generating shadows. These required us to select the area lighting feature for generating shadows from our lighting sources in order to maintain softer shadows with this algorithm. We also created a new stimulus that was similar to the rectangle tested in the transfer phase of Experiment 4, but now having greater height and depth (1 standard unit). Training recommenced 2 weeks later with this improved stimulus set. The first three sessions of testing with the new stimulus set were organized into the same 82-trial organization used at the end of phase 2. At this point, we then added the new rectangular stim-

ulus to their training repertoire. Its eight trials (4 S+/4 S– with four camera angles and four lighting directions randomized) replacing a corresponding number of trials previously testing the sinusoidal surface in each session. These 82-trial sessions testing the five shapes were conducted for 10 days.

5.2. Results

5.2.1. Novel surface shapes

The pigeons generally showed good transfer to most of the novel surface shapes. Shown in Fig. 8 are results for the baseline and five novel shapes, collapsed across perspective, as a function of the mean number of pecks over the first 5 s of presentation. The baseline sinusoid surface continued to support excellent discrimination. Most importantly, the mean number of pecks for each novel shape's positive condition was higher than for its negative condition. A repeated measures ANOVA (Surface Direction \times Shape \times Perspective) using just the results from the five novel shapes confirmed the presence of this significant discrimination transfer as reflected in the main effect of surface direction, $F(1,4) = 85.8$, partial $\eta^2 = .96$. This analysis also revealed a marginally significant interaction between surface direction and shape, $F(4,16) = 2.9$, partial $\eta^2 = .42$, $p = .057$. The source of this interaction can be seen in Fig. 8 as not all of the novel shapes supported equally good transfer, particularly the “bubble” surface. This analysis also revealed a significant interaction between surface direction and camera perspective, $F(3,12) = 9.7$, partial $\eta^2 = .71$. While clear discrimination was present at all the perspective values, it did vary in size among perspectives. The nature of this interaction is more fully explored in the following analyses.

We next looked at discrimination transfer separately for each novel surface shape. The results for each shape were subjected to a repeated measures ANOVA (Surface Direction \times Perspective). Discrimination transfer to the sphere shape was significant, $F(1,4) = 12.4$, partial $\eta^2 = .76$, and its peck rates across the positive and negative conditions did not interact with camera perspective. All five pigeons showed mean higher peck rates to their positive shape. Discrimination transfer to the mound shape was also significant, $F(1,4) = 12.5$, partial $\eta^2 = .76$. All five pigeons showed mean higher peck rates to their positive shape. This analysis found a significant interaction between surface direction and camera perspective, $F(3,12) = 4.9$, partial $\eta^2 = .55$. This interaction was due to the discrimination of this shape being better at the two higher perspectives than the two lower ones. Discrimination transfer to the zigzag shape was also significant, $F(1,4) = 108.3$, partial $\eta^2 = .96$. Surface direction did not interact with camera perspective for this shape, with all five pigeons showing consistently higher peck rates to their positive direction across all camera perspectives. Discrimination transfer to the rectangular shape was also significant, $F(1,4) = 24.5$, partial $\eta^2 = .86$, with again no interaction found between surface direction and camera perspective for the shape. All five pigeons showed mean higher peck rates to their positive direction. Finally, discrimination transfer to the bubble shape was poor and not significantly above chance, $F(1,4) = 1.1$. While

three of the birds showed numerically higher peck rates to their positive direction, none of the birds exhibited the same level of discriminative behavior observed with the other novel shapes.

5.2.2. Novel shape training and lighting transfer

Differential training with the three novel shapes proceeded smoothly. Discrimination started at levels comparable to that observed during transfer and then rapidly and incrementally improved for all three new shapes over the next 10 sessions as they become more familiar. At this point, discrimination leveled off and matched that for the original sinusoidal surface. Over the last twenty sessions of this differential training, discrimination, as measured by the number of pecks in the first 5 s, was excellent and comparable across the surface shapes (sinusoid- S^+ = 13.7, S^- = 2.7; sphere- S^+ = 12.9, S^- = 3.8; mound- S^+ = 15.4, S^- = 3.3; zigzag- S^+ = 15.4, S^- = 6.6). All of these differences were significant as confirmed by paired t -tests. Thus, not surprisingly given the transfer performance, the pigeons easily expanded their repertoire to include this broader variety of shapes.

We next examined how this established discrimination transferred to the two unfamiliar and untrained lighting directions of these shapes. Overall, the pigeons showed excellent transfer to these new stimuli. Using only the number of pecks in the first 5 s of presentation for the three trained previously novel shapes, discrimination combined across the two unfamiliar lighting directions was very good (S^+ = 13.1, S^- = 4.9) and comparable to that observed on baseline trials from the two trained lighting directions (S^+ = 13.6, S^- = 3.1). A repeated measures ANOVA (Surface Direction \times Lighting Direction [combined familiar vs. combined novel] \times Camera perspective) examining only the novel shapes revealed a large main effect of surface direction, $F(1,4) = 19.6$, partial $\eta^2 = .83$, and its interaction with camera perspective, $F(3,12) = 7.7$, partial $\eta^2 = .66$. This interaction with camera perspective was again due to the slightly poorer discrimination supported by the more overhead camera perspectives. Most importantly, however, no significant main effect or interactions were found as related to the familiarity with the lighting directions. The latter confirms the excellent transfer of the novel shape discrimination from the two trained and familiar lighting directions to the two novel lighting directions.

5.2.3. Amended set training

Correction of the absent shadows in the front-lit conditions, the addition of a new shape, and the employment of a better ray tracing algorithm essentially had no effect on the pigeons' established discrimination. Based on just the first three sessions of testing with the new stimulus set, discrimination as measured by the number of pecks in the first 5 s continued to be excellent and comparable across the surface shapes (sinusoid- S^+ = 12.8, S^- = 2.4; sphere- S^+ = 11.9, S^- = 4.0; mound- S^+ = 11.1, S^- = 4.0; zigzag- S^+ = 13.0, S^- = 4.2). All of these differences were found to be significant as determined by paired t -tests. The subsequent introduction of the modified rectangular shape also proved no problem (S^+ = 12.8, S^- = 3.7) as all five pi-

geons significantly discriminated its surface direction during its first session, $t(4) = 3.14$.

The addition of the shadow in the front lit conditions also was no problem for the pigeons. Discrimination across the four light directions proved to be excellent and not significantly different (45° - S^+ = 13.4, S^- = 2.4; 135° - S^+ = 11.6, S^- = 3.5; 225° - S^+ = 12.9, S^- = 2.8; 315° - S^+ = 13.4, S^- = 2.4). To make the most direct comparison possible with the prior front-lit conditions without the shadow, we compared performance from the last three sessions of phase 2 with the first three sessions of this phase. Mean discrimination with the prior unshadowed lighting condition (S^+ = 12.5, S^- = 2.9) was not significantly different from when shadow was properly rendered (S^+ = 13.4, S^- = 2.4), suggesting these occluded shadows were indeed not having much impact.

5.3. Discussion

Overall, this experiment revealed good discrimination transfer to novel surface shapes in several contexts. First, the pigeons showed reliable above-chance transfer to four of the five novel surface shapes upon their initial exposures in the first phase. The pattern of this success strongly suggests it was mediated by the convexity and concavity of the shape as derived from shading of the surface. Second, the pigeons showed very good transfer to these same shapes when the spatial configuration of the shading in these scenes was shifted by 90° in the second phase. Finally, they showed no decrement in discrimination with the amended stimulus set when testing started again.

The most important implication of these results is that they suggest the pigeons were likely seeing the 3D shape of the surfaces in these scenes. When faced with shapes (i.e. mound and zigzag) that had no immediate similarity to the training shape or its associated pattern of shading, the pigeons still differentially discriminated the convex and concave nature of these surfaces upon their first exposure and later during shifts to novel lighting directions. Since the shape and shading of these surfaces were different, the discrimination must have been mediated by the pigeons seeing the directional convexity and concavity of them. For that to happen, they would likely need to see the 3D height and depth implied in the renderings of these shapes.

This experiment also corrected the shadowing of the front-lit stimuli. When this was done, the pigeons continued to discriminate very well. Much as suggested by the earlier results, the absence of this feature was not salient to the birds. One obvious reason why is that it was occluded the vast majority of the time by the shape of the object. Thus, its absence could have easily gone unnoticed (as it did for us). A second, and more speculative, reason is that shadows generally make more important contributions to the spatial configuration of objects in a scene than to shape recognition in humans (Mamassian, Knill, & Kersten, 1998). The same might hold true for pigeons. Examining the contribution of attached and cast shadows and their contribution to shape recognition vs. spatial layout is clearly a question worth exploring in birds.

6. Experiment 5

Whenever an observer moves the shape and patterning of shading information shifts and transforms as the perspective changes. These dynamic changes in lighting are hardly noticed by humans as we move among and around objects. Nevertheless, the systematic regularities within these shifting patterns of light impart important implicit information about the object's shape and our spatial relation to it. The same transparency to changes in illumination would be equally valuable for pigeons and other birds as they walk and rapidly fly through the world. Thus, the motivation behind Experiment 5 was to examine how the pigeons would react to dynamically changing illumination within the present context. Would this confuse their discrimination of surface direction or would they find it as perceptually transparent as we do?

The experiment has two phases in which the pattern of illumination was altered during stimulus presentation. In the first phase, the single lighting source was rendered as moving in a circular path around the surfaces while the camera's perspective remained stationary. This caused the pattern of the light reflected on the surfaces to transform in a manner similar to a sundial. In the second phase, the camera's perspective was rendered as moving in a circular path around the surfaces while the lighting source was held stationary. This caused the pattern of the light reflected from the surfaces to transform as if the observer was moving or flying around the projective shape. For the symmetric shapes, this latter manipulation is identical to moving the light source, but for the asymmetric surface shapes it caused novel perspectives of these shapes to appear and disappear as the camera moved. Using both changes, we investigated how flexibly the pigeons processed surface shape as the pattern of light continually shifted and transformed within a presentation.

A key assumption to examining this issue is whether the pigeons continually process information sequentially during *go/no-go* presentations. Both within the visual and auditory modalities we have found that this is strongly the case for pigeons (Cook & Brooks, 2009; Cook, Kelly, & Katz, 2003; Cook & Roberts, 2007; Murphy & Cook, 2008). In such settings, we have consistently recorded changes in ongoing pecking behavior that are coincident with changes in the discriminative stimuli being presented – indicative of continuous processing. Thus, we investigated here how the current shape-from-shading discrimination was influenced by dynamic changes in lighting and perspective within an ongoing video presentation of the surfaces.

6.1. Methods

6.1.1. Animals and apparatus

The same pigeons and apparatus were used as in Experiment 4.

6.1.2. Procedure

6.1.2.1. Dynamic lighting. In this experiment, a single lighting source dynamically moved in a circular path around

the shape during its 20-s presentation. This path was set at the same distance as the four static lighting positions and included their locations. The dynamic lighting source started at one of two positions (45° or 225°) around the shape. Each frame of a dynamic sequence was rendered in 1.5° increments into an uncompressed 240-frame AVI file as the lighting source moved around the scene. Videos were generated for three different camera perspectives (35°, 65° and 85°) and tested the sinusoid, sphere, zigzag, and mound shapes. The latter shapes were selected so that two symmetrical and two asymmetrical shapes would be included in the test. This resulted in 48 dynamic test stimuli (four surface shapes × two surface polarities × two starting locations × three camera perspectives). These dynamic stimuli were presented within a trial at rate of 11.1 frames per second. Thus within a standard 20-s presentation the lighting source completed 93% of its path around the scene.

These dynamic lighting test stimuli were randomly inserted in sets of 12 non-reinforced probe trials into a session. A different surface shape was probe tested each session at three camera perspectives and at the two starting locations as both dynamic and static presentations. The remaining trials of each session consisted of randomly generated tests of the remaining shapes from different camera perspectives and lighting directions. If the scene parameters of these randomly selected trials matched the testing set, performance on the static negative trials were included in the analysis. Sixteen 78-trial dynamic test sessions were conducted with each shape being tested equally often.

6.1.2.2. Dynamic perspective. In this experiment the camera was dynamically moved in a circular path around the shape during its 20-s presentation. This rotational path maintained the same height, distance, and framing as for the static perspective. The lighting source was located at one of two positions (45° or 225°) around the shape, as measured against the starting location of the camera's perspective. Each frame of a dynamic sequence was rendered in 1.5° increments into an uncompressed 240-frame AVI video as the camera moved around the scene. Videos were generated for three different camera angles (35°, 65° and 85°) and the sinusoid, sphere, zigzag, and mound shapes. This resulted in 48 dynamic test stimuli (four surface shapes × two surface directions × three camera angles × two lighting sources). These dynamic stimuli were presented within a trial at rate of 11.1 frames per sec resulting in the camera traveling 93% of its path around the scene. These dynamic perspective test stimuli were randomly inserted in sets of 12 non-reinforced probe trials into a session. A different surface shape was probe tested each session at three camera perspectives and at the two starting locations as both dynamic and static presentations. The remaining trials of each session consisted with randomly generated tests of the remaining shapes from different camera perspectives and lighting directions. Again, responses from the static negative stimuli of this randomly selected stimulus set were also used in the data analysis. Sixteen 78-trial dynamic test sessions were conducted with each shape being tested equally often.

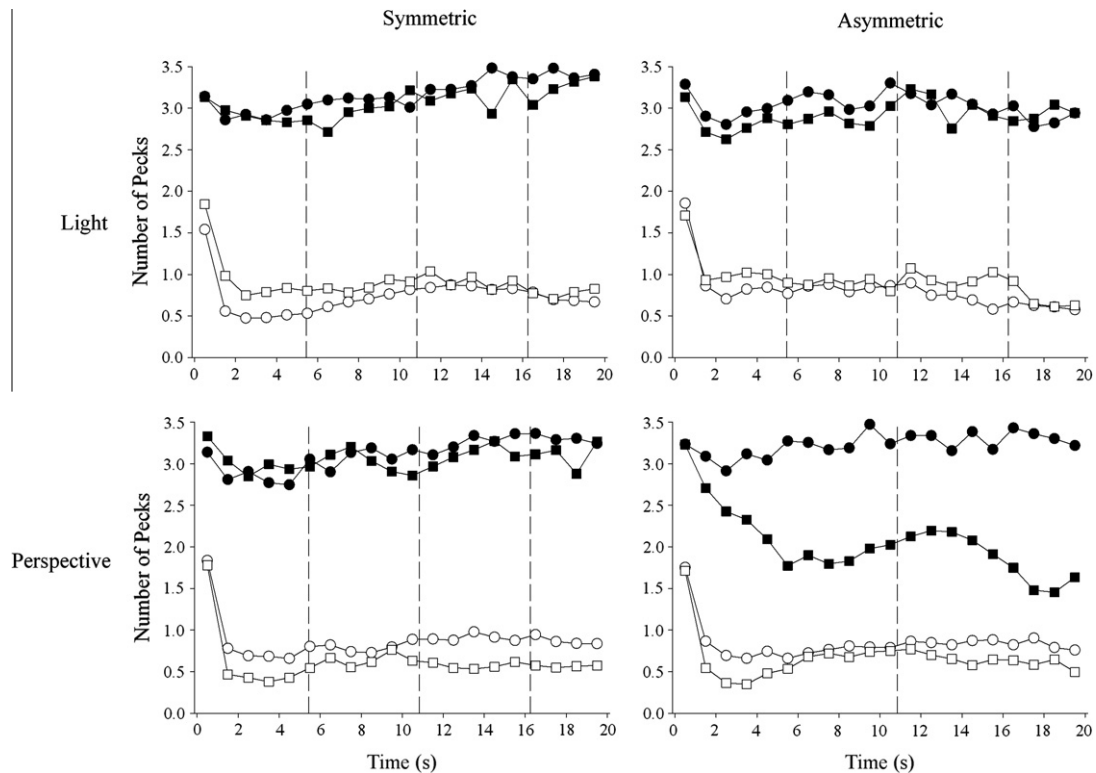


Fig. 9. Mean number of pecks as a function of presentation time in 1-s bins for the dynamic and static S+ and S- conditions tested in Experiment 5 (filled symbol = S+, open symbol = S-; circle = static, square = dynamic). The dashed lines indicate the time that previously reinforced familiar views of static shapes would have been appearing in the dynamic conditions. The top two panels show the results of the dynamic lighting tests, while lower two panels show the results of the dynamic perspective tests.

6.2. Results

6.2.1. Dynamic lighting

Overall the pigeons had few issues in processing the dynamic shifts in lighting during this experiment. Displayed in Fig. 9 are the differences in peck rate for the positive and negative static and dynamic lighting conditions as a function of time in trial in 1-s intervals. The top two panels present these results separately for the two symmetric and two asymmetric surface shapes. As described in Experiment 1, the pigeons quickly discriminated the surface direction of the static displays within a second and maintained this throughout its presentation. When compared to the static condition, discrimination of the dynamic lighting condition was essentially comparable for all surface shapes, despite the latter's continually shifting illumination and resulting changes in shading and shadows.

These curves were subjected to repeated measures ANOVAs (Surface Direction \times Dynamic/Static \times Time) using peck rates in 1-s bins for the symmetric and asymmetric shapes. The ANOVA for the symmetric shapes revealed three significant two-way interactions, besides the obvious main effect of surface direction, $F(1,4) = 16.8$, partial $\eta^2 = .81$. The most important of these was a surface direction \times dynamic/static interaction, $F(1,4) = 49.7$, partial $\eta^2 = .93$. This interaction was due to the slightly better

performance overall with the static condition relative to the dynamic, especially during early portions of the presentation interval. There was also a significant surface direction \times time interaction, $F(19,76) = 1.9$, partial $\eta^2 = .32$, as the relative size of the discrimination changed with time. Finally, there was a dynamic/static \times time interaction, $F(19,76) = 2.6$, partial $\eta^2 = .93$. The outcome of the ANOVA for the asymmetric shapes was simpler. There was only a main effect of surface direction, $F(1,5) = 25.0$, partial $\eta^2 = .86$. Thus, for these shapes the static or dynamic nature of the lighting made no difference.

We also varied the horizontal elevation of the camera in this discrimination. Effects of this manipulation were similar to those reported elsewhere in the article; with increasingly overhead perspectives generally reducing, but not eliminating, the discrimination. We also conducted repeated measures ANOVAs (Surface Direction \times Camera Perspective \times Dynamic/Static) looking at peck rate collapsed across time. There were significant surface direction \times camera perspective interactions for both symmetric, $F(2,8) = 5.7$, partial $\eta^2 = .59$, and asymmetric shapes, $F(2,8) = 19.9$, partial $\eta^2 = .83$, reflecting this progressive decrease in discrimination. These effects were equivalent for both dynamic and static conditions, however, as the ANOVAs do not indicate any significant interactions with this factor.

6.2.2. Dynamic perspective

The lower two panels of Fig. 9 show peck rate for the positive and negative static and dynamic perspective conditions as a function of time in trial (1-s intervals). While the symmetric and asymmetric nature of the shapes produced little difference with dynamic lighting, this characteristic had a larger effect here. As the camera perspective changed, the symmetric shapes showed no impact since they maintained the same shape over time despite the changing illumination. The discrimination of the asymmetric shapes, on the other hand, was affected by this perspective change. As each trial started with a familiar viewpoint, the initial discrimination of these shapes was good. As the camera's viewpoint moved towards perspectives less similar to the training, however, performance declined. As the perspective moved back towards a more familiar viewpoint, discrimination improved and then fell off as less familiar perspectives were displayed. Nevertheless, across all of these perspectives the pigeons' discrimination of the shapes remained significantly above chance. These temporal shifts in discrimination confirm that the pigeons were indeed continually processing the videos during their presentation.

We subjected these curves to repeated measures ANOVAs (Surface Direction \times Dynamic/Static \times Time) using peck rates within the 1-s bins, again separately for the symmetric and asymmetric shapes. For symmetric shapes, this ANOVA revealed a similar set of significant two-way interactions as before, besides the standard main effect of surface direction, $F(1,4) = 16.7$, partial $\eta^2 = .81$. Again, the most important of these was the surface direction by dynamic/static interaction, $F(1,4) = 23.5$, partial $\eta^2 = .85$. This time the interaction was due to the slightly better discrimination overall with the dynamic condition relative to the static, although the absolute size of this effect was small. There was also a significant dynamic/static \times time interaction, $F(19,76) = 1.8$, partial $\eta^2 = .31$. Given the pattern of the curves for the different conditions for the asymmetric shapes in Fig. 9, it is not surprising that there was a significant surface direction \times dynamic/static \times time three-way interaction, $F(19,76) = 4.7$, partial $\eta^2 = .54$. The source of this interaction was clearly due to the generally reduced and changing size of the discrimination with dynamic condition compared to the static condition over time.

To examine the contribution of camera perspective, we conducted similar repeated measures ANOVAs (Surface Direction \times Camera Perspective \times Dynamic/Static) looking at overall peck rate collapsed across time. Like in the first test, there were significant camera perspective \times surface direction interactions for both the symmetric, $F(2,8) = 4.5$, partial $\eta^2 = .53$, and asymmetric shapes, $F(2,8) = 20.5$, partial $\eta^2 = .84$, reflecting this progressive decrease in discrimination with increasing camera perspective. These effects, however, were again equivalent for both dynamic and static conditions, as there were no significant interactions with this factor. Although not visible in the summary provided in Fig. 9, it should be noted that at the 35° camera perspective the pigeons performed the dynamic perspective discrimination at levels that were equivalent to the static condition. It was primarily at the higher camera perspectives that the reduced discrimination of the dynamic condition was most evident.

6.3. Discussion

Experiment 5 revealed the pigeons could readily accommodate to dynamically changing illumination of the surface shapes. Shifts in illumination produced by changes in the location of the lighting source had little impact on their discrimination. This is contrary to the more moderate transfer to novel illumination transfer reported by Young et al. (2001). The discrepancy likely results from the much larger training set with different sources of illumination used here, which may have promoted a more generalizable shape based response to our changes in lighting. Discrimination with changes in the observer's perspective relative to the light source was also generally good. Here, however, the asymmetrical shapes produced temporal changes in discrimination that varied with the familiarity of the current perspective.

Unfamiliar perspectives of the surfaces caused reduced, but still significant, discrimination when in view. During motion of the light source, the pattern of light and dark pixels of the image changed constantly, but in a manner that did not disrupt the discriminative behavior of the pigeons. Instead, the critical response feature is some property that depends on perspective changes, similar to properties critical to shape discrimination. The maintained nature of the discrimination by the pigeons in the face of these dynamic conditions suggests that lighting changes, but not shape changes, are cognitively transparent to them. This transparency in the processing of illumination in comparison to the selective shape-related decrement observed with unfamiliar perspectives is consistent with the hypothesis that the critical feature mediating the pigeons' discrimination is not the shading itself but the shape of the surface derived from the shading.

7. General discussion

Overall, these experiments suggest that the pigeon visual system can perceive shape-based information derived from the patterns of reflected light portrayed from an aerial perspective. Throughout the experiments, the pigeons readily discriminated the direction and shape of surface features transmitted by differential shading information across a wide variety of static and dynamic illumination conditions and did so rapidly (<750 ms). The discrimination generalized to novel perspectives in two axes, changes in the specularity of the reflected light, to novel 3D shapes, and dynamically changing illumination conditions. Thus, whenever challenged with a new organization or configuration of illumination in a scene, the pigeons continued to extract the surface's direction with little difficulty.

The critical question concerns what information the pigeons were using from these different illumination conditions to determine the convexity and concavity of the surfaces. For humans, these stimuli generate a strong impression of depth-based surfaces that allow us to easily see their convexity and concavity. The same seems to be true for the pigeons. Consider first the alternative that they use the 2D patterns of luminance in the stimuli to learn and generalize the discrimination. If so, there are a number

of places where this image-based 2D account fails to account for the present results. For instance, when the camera's perspective was changed in Experiment 2, performance selectively deteriorated for increasingly overhead viewpoints and not for perspectives closely to the horizon. Both of these changes caused changes in the relative patterns of luminance, but only when shape-based information was reduced did the pigeons' discrimination suffer. The same argument would apply to the increasing and decreasing alternations of height in that experiment and to the results using varying specularities in Experiment 3. The transfer to novel shapes in Experiment 4 is also difficult for an image-based mechanism as the novel shapes included extended regions (mound), flat areas (rectangle) and multiple shading gradients (zigzag) which were absent from the original training stimuli. An image-based explanation focused on contour shape would also face great difficulty in this experiment as the contrast contours created by the new shapes differed drastically, including areas of primarily straight horizontal lines (mound, rectangle) and multiple contour form (zigzag). Despite very different resulting patterns of shading, the pigeons were still able to extract the direction of these surfaces and continue to recognize them even when the illumination conditions were inverted and they were lit from novel directions. Finally, in Experiment 5, as the location of the lighting source was moved, altering the luminance patterns in ways not previously experienced by the pigeons, they showed no decline in discrimination. Solely when shape-based information changed with the addition of novel perspectives did the pigeons show any reduction in discrimination. In each of these experiments, the overall patterns of luminance were substantially changed by the experimental manipulations, but the pigeons consistently reacted to the stimuli as if controlled by the convexity of the shape apparent in the image. In fact, the pigeons generally reacted as if the changes in illumination were essentially transparent to them, despite its carrying the critical perceptual information of the scene.

Instead, their discrimination declined when shape-based descriptions of the scenes seemed to offer the best account. Thus, as the camera's perspective moved overhead (Experiment 2), or the shape was increasingly flattened (Experiment 2), or the orientation of an asymmetrical shape was made novel (Experiment 5), did these changes interfere with the pigeons' abilities to discriminate the stimuli. Therefore, we suggest the data are most compatible with the hypothesis that the pigeons were processing the illumination and shading from the scene to obtain 3D information that allowed discrimination of the stimuli based on their shape and depth-based properties. By perceiving depth and shape-from-shading, the pigeons could use the direct abstraction of a shape's convexity and concavity to guide their rapid responding to the displays. In this account, the exact pattern of light and dark is not critical, as long as it perceptually conforms to the structure of a 3D shape. Across the various experiments and many changes in the illumination conditions, only these shape-based features remained invariant. Thus, these results confirm the hint in Cavoto and Cook (2006) that differential shading is likely an important monocular

depth cue in pigeons and suggest the pigeons were likely seeing these illuminated stimuli much as intended – as receding surfaces containing concave or convex shapes that differed in height and depth. If so, they may have shared a similar perception of these stimuli with humans.

We had at least four types of lighting cues present in the scenes that could have contributed to this perception. The first was the relative shading on the surface shapes themselves. Because they carried so much information, they were likely among the most important cues present. Shadows were also present. The exact contribution of the shadows to these shape discriminations is more difficult to isolate. In general the presence and absence of shadows, whether by nature (their relative visibility from front-lit and back-lit scenes) or inadvertently (our initial error with front-lit scenes), did not seem to impact the discrimination at any point. Cavoto and Cook (2006) similarly found that shadows did not influence their discrimination either. Nevertheless, future work will need to more explicitly examine the contributions of shadows to avian visual perception. A third form of lighting was the reflective specular highlights on the shapes. Their variation in Experiment 3 did not dramatically affect the discrimination. Finally, a fourth form of light-derived cues were the contrast contours present in some of the scenes when the foreground and background illumination provided an outline to a shape's contour or spatial boundary, especially at lower camera perspectives, or produced by occlusion of the background. The contours of objects derived from such cues make important contributions to human perception and machine vision. Although we attempted to minimize the presence of such contours in our stimuli, they are quite difficult to eliminate entirely. As a result, they may have made contributions to the birds' perception of the stimuli as well. Working in tandem with shading cues (Cavanagh, 1987; Todd, 2004), they may have helped to create a stronger perception of three-dimensionality and information concerning the direction of the different surfaces.

The relative contributions and priorities of these different lighting cues is beyond the scope of the present investigation and remain to be detailed for this species. While the present experiments indicate that the pigeon visual system, like the human one, can process lighting to extract shape information within a scene without being tied to its specific configuration, the relative salience of these features could differ. Thus, some degree of caution is needed before one says that both species see exactly the same percept or that shape-from-shading is computed using comparable mechanisms. Thus, while both species can clearly rely on shading-transmitted information to discriminate the shape in these displays, further work will be needed to investigate exactly how far these perceptual, functional, and computational similarities may extend. For instance, it will be important in the future to test directly how these different types of lighting cues influence the relative judgment of depth in similar scenes (e.g. Norman et al., 2004).

The results also contribute to the recent debate concerning whether pictures are viewed as representational or not by animals (Bovet & Vauclair, 2000; Spetch & Friedman, 2006; Weisman & Spetch, 2010). The current results seem most consistent with the idea that the pigeons were

seeing the 3D structure represented within these 2D images. We think that several features in the current experiment may have contributed to this outcome. First, the small, efficient visual system of pigeons has been honed by the demands of perceiving the natural world. Thus, when presented with pictures pigeons may bring to bear perceptual mechanisms only evolved for that function and not to make larger visual “inferences” about an image. The rich set of shading-derived scene and depth cues present in the current experiments would be highly compatible with the kind of mechanisms that pigeons would need to navigate and interact with their visual world. Second, it is vital to draw the birds’ attention to these cues by the contingencies of the task. Again, the 3D shading cues provided here were absolutely critical to deriving the direction of the surface shape needed to perform the discrimination and would have made them highly salient and in need of attention. Third, we promoted the global processing of these scenes by recessing the screen back from the touch-screen’s pecking surface, decreasing the visual angle of the stimuli. Given the pigeon’s apparent preference to attend to locally organized information (Cavoto & Cook, 2001), it is important to design both the stimuli and experiments to encourage global processing as strongly as possible. Moving the presentation screen back likely does that. Combining these different techniques may have allowed the pigeons here to react consistently to our rendered representations, much as if they were looking through a window, and react appropriately. Because of the extensive use of pictures and increasing use of video in animal cognition experiments, determining the necessary and sufficient conditions that support representational picture perception in animals is a critical issue.

Given their highly visual nature and mobile lifestyle, it might have been surprising if pigeons were insensitive to the powerful effects of illumination in scene processing. Given the functional importance of object perception and scene recognition to diurnal birds, the processing of these visual features is likely a widespread and general mechanism. Nevertheless, the visual ecology of any species may still influence the relative priority of such cues. For instance, woodland birds that live in the forest canopy may discount such cues because of the reduced levels of ambient light and the considerable amount of visual noise from other objects (e.g., limbs), random shafts of lights, and numerous cast shadows in this habitat. Either way, the challenge for visual science and comparative cognition is to understand how birds rapidly process this kind of information with the limited neural equipment available to them for this function. Quite a number of algorithms have been proposed for trying to understand the contribution of shading information within the domain of artificial intelligence and computer vision (e.g., Horn & Brooks, 1989; Stewart & Langer, 1997; Zhang et al., 1999). If and how these might be implemented into the neural mechanisms known to comprise the avian visual system represents an important area for investigation.

Finally, research with humans has determined that we have a number of working assumptions that are critically involved in the perception of object and shape-from-shading (Liu & Todd, 2004). One of the oldest recognized

assumptions is that our visual system tacitly assumes that lighting is always from overhead (Ramachandran, 1988). Closely related to this is the idea that most scenes and objects are interpreted as if the observer is looking down on them. A third is that we have a bias for seeing most curved surfaces as being convex rather than concave, since this is more common in the visual world (Langer & Bühlhoff, 2000). While these prior assumptions were not tested here, the stimuli were entirely consistent with them. Little work has explored whether the same implicit assumptions are similarly built into the avian mechanisms for scene perception (Hershberger, 1970; Hess, 1950). Given their seamless integration into our ability to perceive scenes, it will be interesting to determine if these are evolutionarily old perceptual “assumptions,” built in to the visual systems of large mobile diurnal organisms since they began to interact with lighted world of objects, or represent more modern adaptations present within only some phylogenetic lineages.

Acknowledgement

This research was supported by Grant #0718804 from the National Science Foundation. E-mail: Robert.Cook@tufts.edu.Home Page: www.pigeon.psy.tufts.edu.

References

- Blinn, J. F. (1977). *Models of light reflection for computer synthesized pictures*. In Paper presented at the Siggraph '77 Proceedings of the 4th annual conference on computer graphics and interactive techniques. New York, New York.
- Bovet, D., & Vauclair, J. (2000). Picture recognition in animals and in humans: A review. *Behavioral Brain Research*, 109, 143–165.
- Cavanagh, P. (1987). Reconstructing the third dimension: Interactions between color, texture, motion, binocular disparity, and shape. *Computer Vision, Graphics, and Image Processing*, 37, 171–195.
- Cavoto, B. R., & Cook, R. G. (2006). The contribution of monocular depth cues to scene perception by pigeons. *Psychological Science*, 17, 628–634.
- Cavoto, K. K., & Cook, R. G. (2001). Cognitive precedence for local information in hierarchical stimulus processing by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 3–16.
- Cook, R. G., Beale, K., & Koban, A. C. (2011). Velocity-based motion categorization by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 175–188.
- Cook, R. G., & Brooks, D. I. (2009). Generalized auditory same-different discrimination by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 108–115.
- Cook, R. G., Kelly, D. M., & Katz, J. S. (2003). Successive two-item same-different discrimination and concept learning by pigeons. *Behavioural Processes*, 62, 125–144.
- Cook, R. G., Levison, D. G., Gillett, S. R., & Blaisdell, A. P. (2005). Capacity and limits of associative memory in pigeons. *Psychonomic Bulletin & Review*, 12, 350–358.
- Cook, R. G., & Roberts, S. (2007). The role of video coherence on object-based motion discriminations by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 287–298.
- Fagot, J., & Cook, R. G. (2006). Evidence for large long-term memory capacities in baboons and pigeons and its implication for learning and the evolution of cognition. *Proceedings of the National Academy of Science*, 103, 17564–17567.
- Gibson, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Hershberger, W. (1970). Attached-shadow orientation perceived as depth by chickens reared in an environment illuminated from below. *Journal of Comparative and Physiological Psychology*, 73, 407–411.

- Hess, E. H. (1950). Development of the chick's responses to light and shade cues of depth. *Journal of Comparative and Physiological Psychology*, 43, 112–122.
- Horn, B. K. P., & Brooks, M. J. (1989). *Shape from shading*. Cambridge, MA: MIT Press.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160, 106.
- Imura, T., & Tomonaga, M. (2003). Perception of depth from shading in infant chimpanzees (*Pan troglodytes*). *Animal Cognition*, 6, 253–258.
- Jenkins, H., & Sainsbury, R. S. (1970). Discrimination learning with the distinctive feature on positive or negative trials. *Attention: Contemporary theory and analysis* (pp. 239–273). New York: Appleton-Century-Crofts.
- Kniss, J., Premoze, S., Hansen, C., Shirley, P., & McPherson, A. (2003). A model for volume lighting and modeling. *IEEE Transactions on Visualization and Computer Graphics*, 150–162.
- Koban, A. C., & Cook, R. G. (2009). Rotational object discrimination by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 250–265.
- Langer, M. S., & Bülthoff, H. H. (2000). A prior for global convexity in local shape from shading.
- Liu, B., & Todd, J. T. (2004). Perceptual biases in the interpretation of 3D shape from shading. *Vision Research*, 44, 2135–2145.
- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*, 2, 288–295.
- Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.
- Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, 53, 137–151.
- Murphy, M. S., & Cook, R. G. (2008). Absolute and relational control of a sequential auditory discrimination by pigeons (*Columba livia*). *Behavioural Processes*, 77, 210–222.
- Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science*, 15, 565.
- Peissig, J. J., Wasserman, E. A., Young, M. E., & Biederman, I. (2002). Learning an object from multiple views enhances its recognition in an orthogonal rotational axis in pigeons. *Vision Research*, 42, 2051–2062.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, 331, 163–166.
- Reid, S., & Spetch, M. L. (1998). Perception of pictorial depth cues by pigeons. *Psychonomic Bulletin & Review*, 5, 698–704.
- Rowland, H. M., Cuthill, I. C., Harvey, I. F., Speed, M. P., & Ruxton, G. D. (2008). Can't tell the caterpillars from the trees: countershading enhances survival in a woodland. *Proceedings of the Royal Society B: Biological Sciences*, 275, 2539–2545.
- Spetch, M. L., & Friedman, A. (2003). Recognizing rotated views of objects: Interpolation versus generalization by humans and pigeons. *Psychonomic Bulletin & Review*, 10, 135–140.
- Spetch, M. L., & Friedman, A. (2006). Pigeons see correspondence between objects and their pictures. *Psychological Science*, 17, 966.
- Stewart, A. J., & Langer, M. S. (1997). Toward accurate recovery of shape from shading under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19, 1020–1025.
- Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research*, 38, 2259–2275.
- Todd, J. T. (2004). The visual perception of 3D shape. *Trends in Cognitive Sciences*, 8, 115–121.
- Tomonaga, M. (1998). Perception of shape from shading in chimpanzees (*Pan troglodytes*) and humans (*Homo sapiens*). *Animal Cognition*, 1, 25–35.
- Tomonaga, M., & Imura, T. (2010). Pacman in the sky with shadows: The effect of cast shadows on the perceptual completion of occluded figures by chimpanzees and humans. *Behavioral and Brain Functions*, 6, 38.
- Wasserman, E. A., & Anderson, P. A. (1974). Differential autoshaping to common and distinctive elements of positive and negative discriminative stimuli. *Journal of the Experimental Analysis of Behavior*, 22, 491–496.
- Weisman, R., & Spetch, M. (2010). Determining when birds perceive correspondence between pictures and objects: A critique. *Comparative Cognition and Behavior Reviews*, 5, 117–131.
- Williams, L. (1978). *Casting curved shadows on curved surfaces*. In Paper presented at the SIGGRAPH '78 Proceedings of the 5th annual conference on computer graphics and interactive techniques. New York, New York.
- Young, M. E., Peissig, J. J., Wasserman, E. A., & Biederman, I. (2001). Discrimination of geons by pigeons: The effects of variations in surface depiction. *Animal Learning & Behavior*, 29, 97–106.
- Zhang, R., Tsai, P. S., Cryer, J. E., & Shah, M. (1999). Shape-from-shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 690–706.