

# The effects of smooth occlusions and directions of illumination on the visual perception of 3-D shape from shading

Eric J. L. Egan

Department of Psychology, the Ohio State University,  
Columbus, OH, USA



James T. Todd

Department of Psychology, the Ohio State University,  
Columbus, OH, USA



**Human observers made local orientation judgments of smoothly shaded surfaces illuminated from different directions by large area lights, both with and without visible smooth occlusion contours. Test–retest correlations between the first and second halves of the experiment revealed that observers’ judgments were highly reliable, with a residual error of only 2%. Over 88% of the variance between observers’ judgments and the simulated objects could be accounted for by an affine correlation, but there was also a systematic nonaffine component that accounted for approximately 10% of the perceptual error. The presence or absence of visible smooth occlusion contours had a negligible effect on performance, but there was a small effect of the illumination direction, such that the response surfaces were sheared slightly toward the light source. These shearing effects were much smaller, however, than the effects produced by changes in illumination on the overall pattern of luminance or luminance gradients. Implications of these results for current models of estimating 3-D shape from shading are considered.**

a visible surface toward the point of observation is influenced by three factors: (a) the surface geometry, (b) the pattern of illumination, and (c) the manner in which the surface material interacts with light.

It has been well documented in the literature that observers’ judgments of 3-D shape from shading do not remain perfectly constant over changes in the pattern of illumination or surface material properties (Caniard & Fleming, 2007; Christou & Koenderink, 1997; Curran & Johnston, 1996; Khang, Koenderink, & Kappers, 2007; Koenderink, van Doorn, Christou, & Lappin, 1996a, 1996b; Mingolla & Todd, 1986; Mooney & Anderson, 2014; Nefs, Koenderink, & Kappers, 2005; Todd, Koenderink, van Doorn, & Kappers, 1996; Todd & Mingolla, 1983; Wijntjes, Doerschner, Kucukoglu, & Pont, 2012). However, these violations of constancy are typically quite small relative to the overall changes in the patterns of shading on which they are based. A particularly elegant demonstration of this was reported by Christou and Koenderink (1997; see also Koenderink et al., 1996a, 1996b). Observers in this study performed a local attitude adjustment task at numerous probe points on a sphere under varying conditions of illumination. Subsequent analyses were then performed on the data to compute a smooth surface for each condition that was maximally consistent with the overall pattern of adjustments. The results showed clearly that the apparent shapes of the objects were sheared slightly toward the direction of illumination. An additional analysis was performed to compute a best fitting surface from the local luminance gradients within the stimulus displays. These surfaces were also sheared toward the direction of illumination, but the magnitudes of these deformations were much larger than those obtained from the observers’ judgments.

## Introduction

One of the most perplexing aspects of human perception involves the ability of observers to determine the 3-D shapes of objects from patterns of image shading. In an effort to explain how this is possible, researchers have developed numerous algorithms for computing shape from shading (see Zhang et al., 1999, for a review), but the performance of these models has been consistently disappointing relative to that of human observers. Patterns of shading are inherently difficult to interpret because the light that reflects from

Citation: Egan, E. J. L., & Todd, J. T. (2015). The effects of smooth occlusion and directions of illumination on the visual perception of 3-D shape from shading. *Journal of Vision*, 15(2):24, 1–11, <http://www.journalofvision.org/content/15/2/24>, doi:10.1167/15.2.24.

Christou and Koenderink (1997) speculated that the relatively high degree of constancy exhibited by the observers was most likely due to information provided by visible smooth occlusion contours. This appears at first blush to be a reasonable suggestion. Previous research has shown that smooth occlusion contours can provide a powerful source of information for the analysis of 3-D shape from shading. The surface normals along smooth occlusion contours are always perpendicular to the line of sight, which can provide a critical boundary condition for computational analyses (Ikeuchi & Horn, 1981). These contours also provide information about the surface curvature in their immediate local neighborhoods, because the sign of surface curvature in a direction perpendicular to an attached smooth occlusion contour must always be convex (Koenderink, 1984; Koenderink & van Doorn, 1982). Under conditions of homogeneous illumination, the local luminance maxima along smooth occlusion boundaries provide additional information about the tilt of the illumination direction (Todd & Reichel, 1989). There is even some empirical evidence, obtained by Koenderink et al. (1996b), that observers can make remarkably accurate judgments about surface relief from the silhouette of a familiar object presented in isolation without any smooth variations in shading.

Although it is reasonable to suspect that smooth occlusion contours may play an important role in the perception of shape from shading to facilitate constancy over changes in the pattern of illumination or surface material properties, there are surprisingly few data to support that hypothesis. Virtually all previous studies that have investigated shape constancy from shading have included visible smooth occlusion contours in all conditions. The one exception is a recent article by Todd, Egan, and Phillips (2014) that found little or no effect on performance when surfaces were presented at an orientation for which there were no visible self-occlusions. The research reported in the present article was designed to further examine this issue. The methodology employed was quite similar to those used by Christou and Koenderink (1997) and Koenderink et al. (1996a, 1996b), but the stimuli were presented both with and without visible smooth occlusion boundaries.

## Methods

### Participants

Five observers participated in the experiment: both authors and three other observers who were naïve as to the issues being investigated. All of the observers had normal or corrected-to-normal visual acuity, and all

wore an eye patch to eliminate conflicting flatness cues from binocular vision.

### Apparatus

The experiment was conducted using a Dell Precision 1650 PC with an NVIDIA Quadro 600 graphics card. The stimulus images were presented on a 10-bit, 28-in. gamma-corrected LCD with a spatial resolution of  $2560 \times 1440$  pixels. The images were presented within a 32.5-cm  $\times$  32.5-cm region ( $1024 \times 1024$  pixels) of the display screen, which subtended  $18.5^\circ \times 18.5^\circ$  of visual angle when viewed at a distance of 100 cm.

### Stimuli

The stimuli were all generated from one of six possible scenes, each of which included a randomly deformed object inside a cube-shaped room. The x, y, and z dimensions of the room were all 305 cm. The objects were generated by adding a series of sinusoidal perturbations to a sphere at random orientations, and they were all approximately 90 cm in diameter. The objects were positioned in the room just slightly in front of the center of one wall, and the camera was positioned in the center of the opposite wall. The scenes were illuminated by one of six possible area lights located on the same wall as the camera, and they all had an area of 22,040 cm<sup>2</sup>. The scenes included three different objects, and each of these objects could be illuminated from one of two different directions. The images were created using the Finalrender 3.5 renderer by Cebas. A global illumination model was employed with three bounces of interreflection, and all of the surfaces had Lambertian reflectance functions. The resulting images could be presented in full or within a convex aperture that masked all of the occlusion contours on the depicted object. The masks were designed to reveal as much of the surface area as possible but without any concavities that could potentially provide information about the sign of surface curvature. The masked and unmasked stimuli are shown in Figure 1A and B, the shapes and positions of the area lights are shown in Figure 1C, and the patterns of isointensity contours for each of the images are shown in Figure 1D.

### Procedure

The task on each trial was to adjust the slant and tilt of a circular gauge figure centered at a given probe point so that it appeared to be within the tangent plane of the surface at that point (see Koenderink, van Doorn

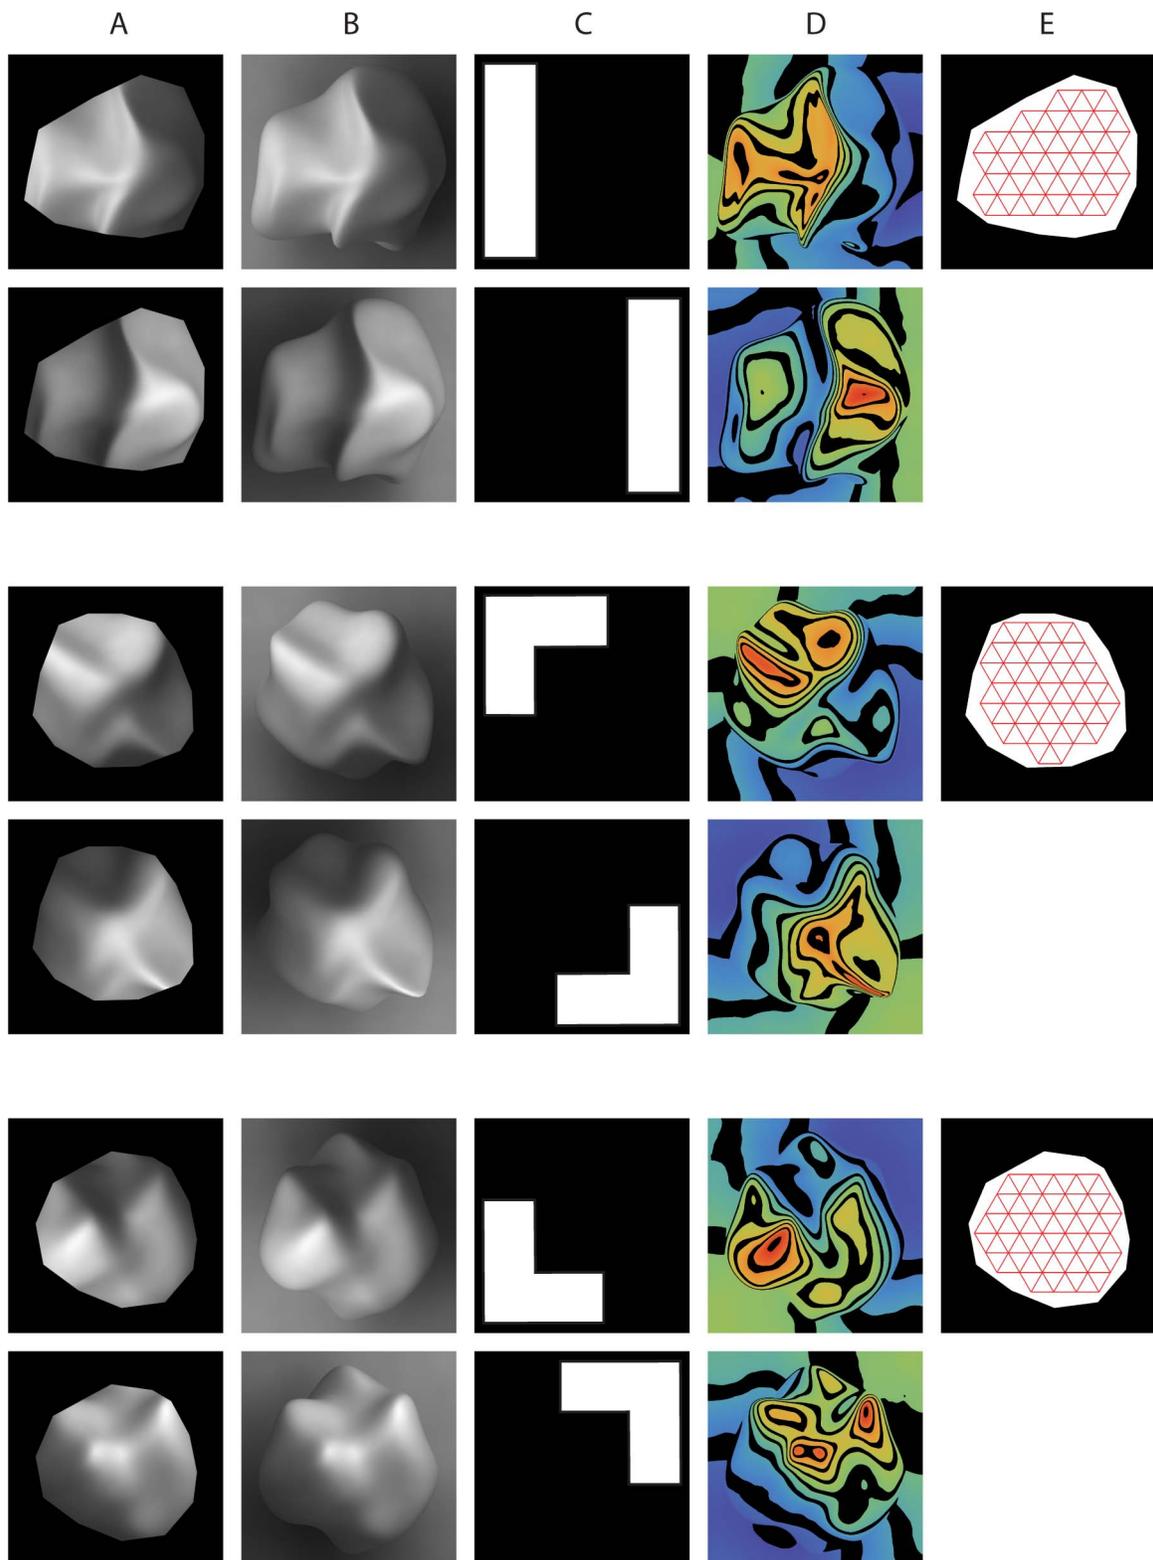


Figure 1. (A) Stimuli from the masked condition. (B) Stimuli from the unmasked condition. (C) The shapes and positions of the area lights with which the objects were illuminated. (D) The patterns of isointensity contours for the unmasked displays. (E) The pattern of probe points used for each object.

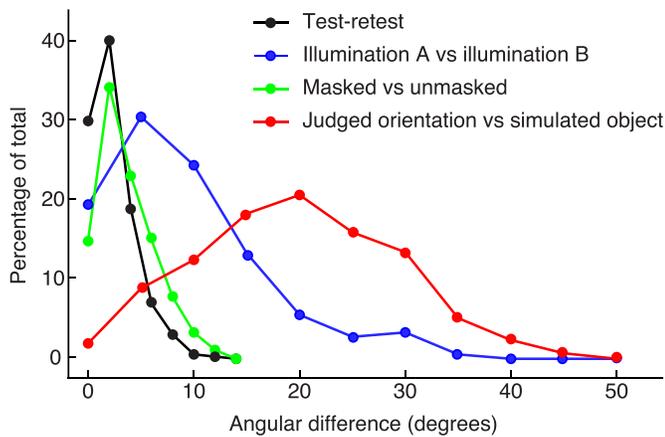


Figure 2. The distribution of test–retest differences between the first and second halves of the experiment (black); the distribution of differences between corresponding probe points in the masked and unmasked conditions (green); the distribution of differences between the corresponding probe points with different directions of illumination (blue); and the distribution of differences between the simulated object and the average of observers’ settings for each probe point in each condition (red).

& Kappers, 1992, 1995). Slant is defined in this context as the angle between the surface normal and the line of sight, whereas tilt is the direction of the surface depth gradient within the frontoparallel plane. The gauge figure simulated a small circle in 3-D space with a radius of 18 pixels and a perpendicular line at its center with a length of 18 pixels. These appeared in the image as a red ellipse with a small line along the minor axis, whose length and orientation could be manipulated using a handheld mouse. When adjusted appropriately, all of the observers were able to perceive this configuration as a circle oriented in depth in the tangent plane with a line perpendicular to it in the direction of the surface normal. Observers typically perform these judgments quite rapidly, at a rate of 8 to 10 per min.

Observers made judgments at 60 probe points for each of the six stimuli. These probe points were located at the vertices of a triangular mesh consisting of 42 faces that evenly covered the majority of each surface (see Figure 1E). The probe points were in identical locations across different illumination conditions for a given surface but differed between objects.

An experimental session was divided into six blocks. Within a given block, observers made settings for all of the different probe points for one of the possible stimulus images. A single experimental session included two blocks for each of the three possible objects with one of its possible directions of illumination. The remaining directions of illumination were presented in different sessions, to avoid any possible interactions between them on observers’ judgments. The entire

experiment comprised eight experimental sessions on separate days. In the first four sessions, all of the stimuli had masked occlusion boundaries (see Figure 1); in the final four sessions, they all had visible occlusion boundaries. This ordering was employed so that observers’ judgments of the masked stimuli could not be influenced by prior knowledge from the unmasked conditions. To summarize briefly, observers made four separate judgments for each of the 60 possible probe points in each condition (3 objects  $\times$  2 illuminations  $\times$  2 types of mask), and the different possible illuminations and masks for a given object were never presented together within the same experimental session. Thus each observer made a total of 2,880 settings over eight different days.

## Results

We began our analysis by measuring the consistency of observers’ judgments over multiple experimental sessions. In order to assess test–retest reliability, we averaged the judgments over all five observers and calculated the angular difference between their settings in the first and second halves of the experiment for each probe point in each condition. The black curve in Figure 2 shows the distribution of these test–retest differences. Note that the angular differences are all tightly clustered with a mean of only 3.3°, which provides a good estimate of the measurement error in this experiment for evaluating other findings.

The red curve in Figure 2 shows the distribution of differences between the simulated object and the average observer settings for each probe point in each condition. The mean of this distribution was 22.6°—almost 7 times as large as the test–retest differences—indicating that there were systematic biases in observers’ judgments. The blue curve in Figure 2 shows the distribution of differences between the average settings at corresponding probe points in each of the two illumination conditions for a given object. The mean of that distribution was 11.5°, which is approximately 3 times as large as the measurement error. This finding is consistent with many previous investigations that have shown that observers’ judgments do not remain perfectly constant over changes in the pattern illumination. Finally, the green curve in Figure 2 shows the distribution of differences between the average settings at corresponding probe points in the masked and unmasked conditions. Note that the mean of that distribution was 4.6°, only slightly larger than the test–retest measurement error. This finding provides strong evidence that the presence of visible smooth occlusion contours in the unmasked displays had relatively little influence on performance. In order to perform a

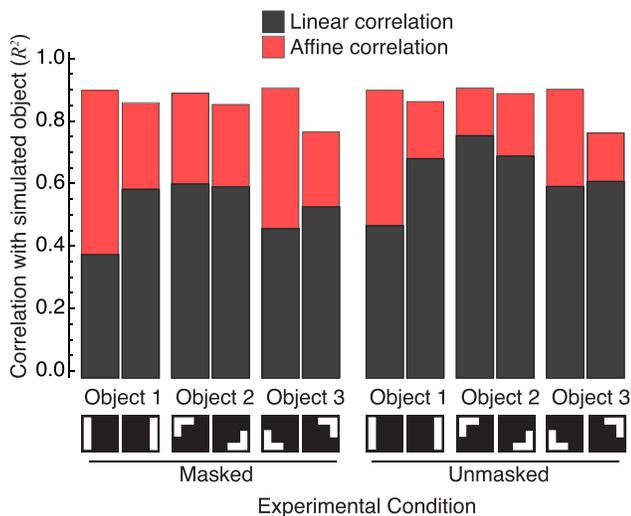


Figure 3. The linear and affine correlations between the simulated object and the response surface computed from the average of observers' settings in each condition. The small square panels depict the patterns of illumination.

statistical comparison of these distributions, we first computed a log transform of the data to make them approximately normal, and we then performed  $t$  tests on all six possible pairwise combinations. The results revealed that all of these comparisons were statistically significant,  $p < 0.01$ , including the small difference between the test–retest distribution and the masked–unmasked distribution.

In an effort to understand the global pattern of the local orientation settings, we employed an analysis originally developed by Koenderink et al. (1992, 1995), using MATLAB code that has been published by Wijntjes (2012). Each orientation setting in this task can be interpreted as a depth gradient, which can be used to estimate the apparent depth difference between adjacent probe points. All of the depth differences in a given condition define an overdetermined system of linear equations. The least-squares solution to these equations defines a best fitting surface that is maximally consistent with the overall pattern of adjustments (see Wijntjes, 2012). To estimate the reliability of this analysis, we computed a best fitting surface from the first two responses in each condition averaged over observers and correlated those surfaces with the best fitting surfaces from the second two responses. The average coefficient of determination ( $R^2$ ) in that case was 0.98, which confirms that the overall pattern of responses is highly reliable.

Figure 3 shows the linear and affine correlations between the simulated objects and the best fitting surfaces from the average responses over all five observers for each of the 12 conditions. The linear correlations show how well the relative apparent depth

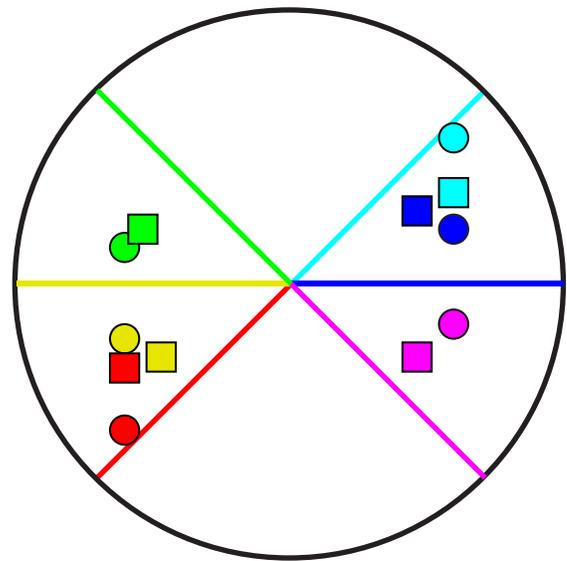


Figure 4. A polar plot of the directions and magnitudes of shear in each condition relative to the average shear for each individual object. The squares and circles represent the masked and unmasked conditions, respectively. The different directions of illumination are coded by color, and the black circle represents a shear magnitude of 0.075.

( $Z'$ ) at each probe point can be predicted by the relative depth ( $Z$ ) of that point on the simulated object. This accounts for approximately 60% of the variance in observers' judgments. The affine correlations show how well the relative apparent depth ( $Z'$ ) at each probe point can be predicted by a linear combination of all three position coordinates of the probe point on the simulated object ( $Z' = aX + bY + cZ$ ). This accounts for approximately 89% of the variance in observers' judgments.

It is important to recognize that all three parameters of the affine correlation reveal specific patterns of distortion between the best fitting response surfaces and the simulated object. The  $c$  parameter, which is equivalent to the slope of the linear correlation, reveals the relative magnitude of apparent depth relative to the simulated object. The average value of that parameter over the 12 conditions of the present experiment was 0.42, which indicates that the judged patterns of relief were compressed in depth by an average of 58%. The  $a$  and  $b$  parameters of the affine correlations reveal shearing transformations between the best fitting response surfaces and the simulated objects. These are best visualized by converting those parameters to polar coordinates so that the directions and magnitudes of shear are represented by the angular and radial components, respectively. Figure 4 shows a polar plot of the apparent shear for each condition relative to the average shear for each individual object. Note that the directions of these shearing transformations were approximately aligned with the directions of illumina-

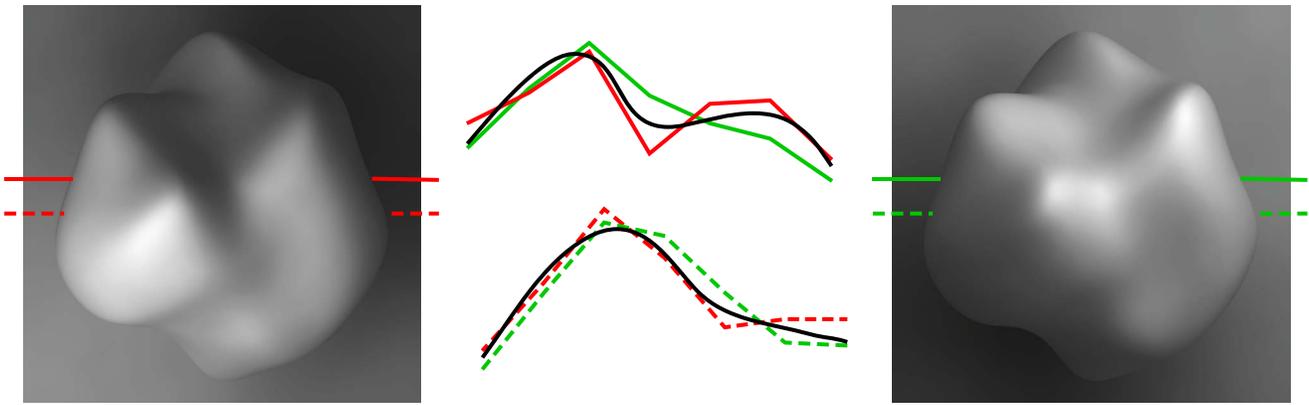


Figure 5. Two scan lines are marked for the images of Object 3 with different directions of illumination. The black curves in the center show the depicted depth profiles along those scan lines. The red and green curves show the computed depth profiles from the best fitting response surfaces, which have been transformed to eliminate any systematic errors due to shearing or depth scaling. Thus, all remaining differences from the simulated object are due to nonaffine perceptual distortions.

tion, as has been shown previously by Christou and Koenderink (1997) and Koenderink et al. (1996a, 1996b). An analysis of linear regression revealed that the tilts of the directions of illumination account for 69% of the variance in the directions of apparent shear among the different possible conditions.

There is another interesting aspect of these data that deserves to be highlighted. From the test–retest correlations of the best fitting surfaces, the residual variance was only 2%. From the affine correlations between the best fitting surfaces and the simulated object, the residual variance was 12%—6 times as large as the measurement error in this experiment. This provides strong evidence that there was a systematic nonaffine component to the distortions of apparent shape. Figure 5 shows the two images of Object 3, which was the one that produced the lowest affine correlations in this experiment; two horizontal scan lines on each image are marked by red and green lines. The graphs in Figure 5 show the depth profiles of each scan line (black curves) and the depth profiles obtained from the observers' judgments (red and green curves), which have been transformed by the best fitting parameters from the affine correlation. Thus, any deviations of the red and green curves relative to the depth profile of the simulated object are due to nonaffine distortions. Note in particular how illumination from the left casts a shadow in the center of the object that accentuates the apparent curvature of the small concavity along the upper scan line, but how the apparent curvature in this region is attenuated by illumination from the right. Similarly, for the lower scan line, illumination from the left causes an inflection point on the surface to appear as a concavity.

In addition to analyzing the average responses over all observers, we also performed similar analyses on the best fitting surfaces obtained from the judgments of individual observers. Figure 6 shows the linear and

affine correlations between the best fitting response surface and the simulated object for each individual in each condition. Figure 7 shows the average linear and affine correlations between each possible pair of observers. In order to assess the measurement error for individual pairs of observers, we computed a best fitting surface from the first two responses of both observers in a given condition and correlated that with the best fitting surface from the second two responses of those observers. The average coefficient of determination over each pair of observers and each condition was 0.98. Because the residual variance from the affine correlations between observers was consistently much larger than the residual variance from the test–retest correlations, these findings indicate that there was also a nonaffine component of the between-observer differences.

Additional analyses were performed to compare observers' performance with some possible heuristics for computing shape from shading that have been proposed in the literature. One such heuristic is the “dark-is-deep” rule that was originally suggested by Langer and Zucker (1994) for diffuse light fields (see also Sun & Schofield, 2012). If observers employed this heuristic, then the relative apparent depths among different local regions should be proportional to the relative luminance in those regions. This clearly did not occur, however. Analyses of linear regression in each condition revealed that the variations in image intensity account for only 7% of the variance in the relative apparent depths among the different possible probe points.

Another possible heuristic is the “linear shape from shading” model proposed by Pentland (1989) for oblique directions of illumination (see also Sun & Schofield, 2012). If observers employed that heuristic, the relative apparent slants among different local regions should be proportional to the relative lumi-

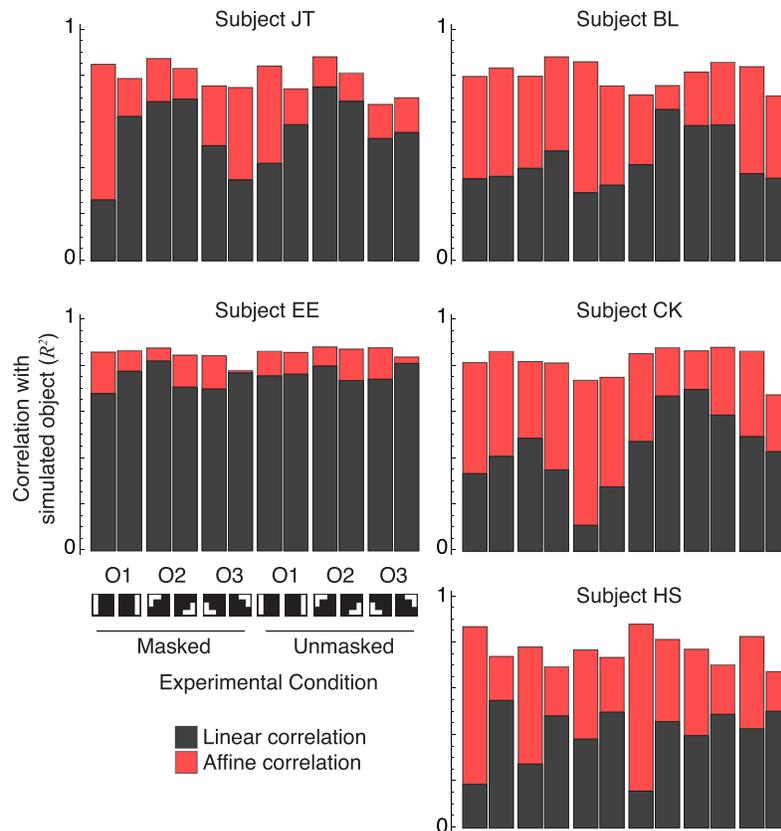


Figure 6. The linear and affine correlations between the simulated object and the response surface computed from the average of observers’ settings in each condition. The small square panels depict the patterns of illumination.

nance in those regions. However, analyses of linear regression again revealed that variations in image intensity account for only 8% of the variance in the relative apparent slants among the different possible probe points.

A third possible heuristic is the “regression to image luminance” model originally proposed by Christou and Koenderink (1997) and Koenderink et al. (1996a, 1996b). This model assumes that observers adjust the

gauge figure so that its slant and tilt are consistent with the magnitude and direction of the decreasing local luminance gradient. It is important to keep in mind that local luminance gradients are primarily influenced by surface curvature and the direction of illumination rather than surface slant, but this is perhaps a plausible strategy for obtaining reliable (albeit inaccurate) slant estimates. In order to test this hypothesis, we used a derivative filter in Mathematica to measure the luminance gradients at the different probe points in each condition. Unit surface normals were then computed from the gradients, and best fitting surfaces were computed using the same analysis that was used for the observers’ judgments. We then correlated the best fitting surfaces obtained from the luminance gradients with those obtained from the gauge figure settings. The results revealed that they were largely independent of one another (i.e.,  $R^2 = 0.02$ ).

The primary goal of the present experiment was to measure the constancy of perceived shape over variations in the pattern of illumination both with and without visible smooth occlusion contours. The results presented thus far have demonstrated that the differences between the two masking conditions were negligible. For a given object and illumination direction, the linear correlations between the best fitting

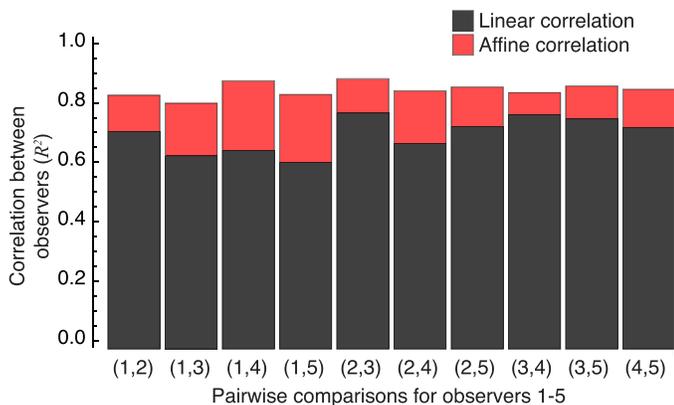


Figure 7. The linear and affine correlations between the best fitting response surfaces for each possible pair of observers averaged over conditions.

surfaces with and without visible smooth occlusion contours revealed an average  $R^2$  of 0.95. The results also demonstrate that the apparent shapes of the depicted objects were systematically distorted by changes in the illumination direction. For a given object and masking condition, the linear and affine correlations between the best fitting surfaces for different directions of illumination revealed average  $R^2$ s of 0.79 and 0.89, respectively.

How should we assess the magnitudes of these failures of constancy over changes in illumination? One possible way to address this issue is to compare them to the changes in image structure. For example, analyses of linear regression between the patterns of image intensities across different illumination directions revealed an average  $R^2$  of 0.49, which is much smaller than correlations between observers' judgments across the same variations in illumination. This is also true for the best fitting surfaces obtained from the luminance gradients. The correlations of those surfaces across different directions of illumination were again quite small (i.e., the average  $R^2$  was 0.10).

Additional analyses were performed to examine the relative stability over changes in illumination for a wide variety of local measures of the luminance field  $L = f(x, y)$ . These included the first and second spatial derivatives of luminance ( $f_x, f_y, f_{xx}, f_{yy}, f_{xy}$ ) and several other measures derived from those, including slant, tilt, mean curvature, Gaussian curvature, curvedness, and shape index (see Koenderink, 1990). These measures were computed for each of the possible probe points in each condition, and a regression analysis was performed to compare their values across the different directions of illumination for each object. None of these correlations produced an  $R^2$  greater than 0.3. These findings suggest that the distortions of apparent shape due to changes in illumination are relatively small when compared to the changes in the zero-order, first-order, and second-order differential structures of the luminance patterns. That is to say, observers are somehow able to achieve a substantial degree of constancy over changes in illumination, although it is by no means perfect.

## Discussion

One of the basic findings of the present investigation is that observers' perceptions of 3-D shape from shading are sheared slightly toward the direction of illumination. This phenomenon was first discovered by Koenderink et al. (1996a, 1996b) and Christou and Koenderink (1997). They argued that this effect is due to a tendency of observers to adjust the gauge figure so that its slant and tilt are consistent with the magnitude

and direction of the decreasing local luminance gradient, which they referred to as a regression to image luminance. However, they also noted that the distortions of apparent shape revealed by the observers' settings were much smaller than what would be expected from that response strategy. They speculated that the residual constancy was most likely due to information provided by visible smooth occlusion contours. The present investigation was designed specifically to test that suggestion, and the results show clearly that the absence of smooth occlusions in the masked conditions had a negligible impact on the overall pattern of performance.

It is important to note that there are some discrepancies in the literature on this point. One study by Koenderink et al. (1996b) has shown that observers can make remarkably accurate judgments about surface relief from the silhouette of a familiar object (i.e., a human torso) presented in isolation without any smooth variations in shading. A similar demonstration was also reported by Mamassian and Kersten (1996) using the silhouette of an unfamiliar croissant-shaped object. However, the observers in that experiment were prompted by the instructions to visualize one side of the object as being closer in depth than the other, and it is reasonable to question whether the results would have been the same in the absence of those instructions.

Although it is quite likely that observers obtain useful information from smooth occlusion contours (see Todorović, 2014), it does not necessarily follow that they are unable to successfully analyze patterns of shading in their absence. The first experiment to address this issue was performed by Todd and Reichel (1989). They asked observers to make relative local depth judgments for small patches of shading taken from a larger image of a 3-D surface. Performance was highest when the stimuli contained a visible smooth occlusion contour regardless of patch size. For the smallest patches without smooth occlusions, the accuracy of the observers' judgments was no greater than chance, but performance improved significantly as the patch size was increased. A follow-up experiment by Reichel and Todd (1990) examined the interaction of visible smooth occlusions with the strong bias of most observers to interpret shaded images of slanted surfaces so that depth appears to increase with height in the visual field. Their results showed clearly that visible smooth occlusions can override this bias, but only in their immediate local neighborhoods. When this occurs, observers' relative depth judgments can be globally inconsistent, much like the appearance of an impossible figure.

Another important factor in the present experiment that may have facilitated observers' judgments in the masked conditions was the model of illumination used to render the displays. The stimuli were created using a

global illumination algorithm that simulated the effects of light attenuation, cast shadows and surface interreflections in a physically realistic manner. These effects have typically been ignored in most studies of the perception of 3-D shape from shading, but they can provide useful information in appropriate contexts. For example, when objects are illuminated with extended light sources, concave regions receive a smaller amount of illumination than convex regions, which is often referred to as vignetting. Langer and Bühlhoff (2000) have shown that this provides sufficient information to accurately distinguish concavities from convexities in regions that are far removed from any visible smooth occlusions, and we suspect this may also have provided a useful source of information in the present experiment.

Given that changes in the pattern of illumination can produce large distortions in the overall pattern of luminance, how is it possible that their effect on apparent shape is so much smaller? There are two possibilities that have been suggested in the literature. One is that available information about the pattern of illumination is somehow taken into account by the mechanisms for computing 3-D shape from shading (Pentland, 1982). Note in Figure 1, for example, that the approximate tilt of the illumination direction is easily determined from the pattern of isointensity contours. There is some empirical evidence to indicate that judgments of shape and illumination may be largely independent of one another (e.g., Khang et al., 2007; Mingolla & Todd, 1986), but these studies have been limited to ellipsoid surfaces with no internal cast shadows or interreflections, so they may not generalize to more natural scenes such as those used in the present experiment.

Another possibility is that observers make use of higher order spatial derivatives of the luminance field  $L = f(x, y)$ . Kunsberg and Zucker (2013) have argued that these higher order structures are more stable over changes in illumination than are patterns of luminance or luminance gradients. Our analyses of seven local curvature measures did not confirm this hypothesis. These measures were  $f_{xx}$ ,  $f_{yy}$ ,  $f_{xy}$ , mean curvature, Gaussian curvature, curvedness, and shape index, which are commonly used measures in differential geometry. However, this does not preclude the possibility that some other set of measures exist that remain relatively invariant over changes in illumination direction. One interesting aspect of Kunsberg and Zucker's model is that it is not purely local. It proposes that information can propagate across different regions of a surface in order to reduce the inherent ambiguity of local shape estimates. We believe this is a promising suggestion that is deserving of future research.

The visual perception of shape from shading has been an active area of research for over three decades, but an adequate theoretical explanation of this phenomenon has proven to be quite elusive. Using the techniques

developed by Koenderink and colleagues for computing a best fitting surface from an overall pattern of local orientation or relative depth judgments, we are now able to predict how apparent shapes can be systematically distorted and how they are influenced by a variety of contextual factors. To a first approximation, observers' judgments are consistent with the bas-relief ambiguity originally described by Belhumeur, Kriegman, and Yuille (1999; see also Koenderink, van Doorn, Kappers, & Todd, 2001). That is to say, the apparent magnitudes of relief for real or simulated objects are often systematically compressed in depth, and their apparent shapes can also be distorted by shearing transformations in the direction of illumination or the overall direction of object slant (see Koenderink et al., 2001; Todd et al., 2014). These two effects account for most of the systematic errors in observers' judgments. Other research has shown, however, that there are also some systematic nonaffine distortions of perceived shape (Caniard & Fleming, 2007; Nefs et al., 2005; Todd et al., 2014), which are less well understood.

Perhaps the most difficult aspect of the perception of shape from shading to explain theoretically is the ability of observers to cope with changes in the manner or direction of illumination, including the effects of cast shadows and surface interreflections, and variations in surface material properties, including transparency, translucency, and specular highlights. These phenomena can produce large changes in the overall pattern of luminance or luminance gradients, yet their effects on observers' perceptions are often relatively small in comparison. We do not wish to suggest, however, that this is inevitably the case. It is clear that there are some situations where our ability to disentangle the effects of shape, illumination, and material properties breaks down. A particularly good example of this has recently been provided by Mooney and Anderson (2014). Using natural light fields, they demonstrated that there are some conditions in which it is difficult to distinguish the specular and diffuse components of shading, and that this can cause specular highlights to be perceptually misinterpreted as variations in local 3-D shape.

*Keywords:* 3-D surface and shape perception, shading, surfaces

## Acknowledgments

This research was supported by a grant from the National Science Foundation (BCS-0962119).

Commercial relationships: none.

Corresponding author: James T. Todd.

Email: todd.44@osu.edu.

Address: Department of Psychology, the Ohio State University, Columbus, OH, USA.

## References

- Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision*, 35(1), 33–44.
- Caniard, F., & Fleming, R. (2007). Distortion in 3D shape estimation with changes in illumination. In C. Wallraven & V. Sunstedt (Eds.), *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization*, 253, 99–105.
- Christou, C. G., & Koenderink, J. J. (1997). Light source dependence in shape from shading. *Vision Research*, 37(11), 1441–1449.
- Curran, W., & Johnston, I. A. (1996). The effect of illuminant position on perceived curvature. *Vision Research*, 36(10), 1399–1410.
- Ikeuchi, K., & Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1–3), 141–184.
- Khang, B.-G., Koenderink, J. J., & Kappers, A. M. L. (2007). Shape from shading from images rendered with various surface types and light fields. *Perception*, 36(8), 1191–1213.
- Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, 13, 321–330.
- Koenderink, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.
- Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, 11, 129–137.
- Koenderink, J. J., van Doorn, A. J., Christou, C., & Lappin, J. S. (1996a). Shape constancy in pictorial relief. *Perception*, 25, 155–164.
- Koenderink, J. J., van Doorn, A. J., Christou, C., & Lappin, J. S. (1996b). Perturbation study of shading in pictures. *Perception*, 25, 1009–1026.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, 52, 487–496.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1995). Depth relief. *Perception*, 24(1), 115–126.
- Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2001). Ambiguity and the “mental eye” in pictorial relief. *Perception*, 30(4), 431–448.
- Kunsberg, B., & Zucker, S. (2013). Characterizing ambiguity in light source invariant shape from shading. Retrieved from <http://arxiv.org/abs/1306.5480>
- Langer, M. S., & Bülthoff, H. H. (2000). Depth discrimination from shading under diffuse lighting. *Perception*, 29, 649–660.
- Langer, M. S., & Zucker, S. W. (1994). Shape-from-shading on a cloudy day. *Journal of the Optical Society of America A*, 11(2), 467–478.
- Mamassian, P., & Kersten, D. (1996). Illumination, shading and the perception of local orientation. *Vision Research*, 36, 2351–2367.
- Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, 53, 137–151.
- Mooney, S. W. J., & Anderson, B. L. (2014). Specular image structure modulates the perception of three-dimensional shape. *Current Biology*, 24, 2737–2742.
- Nefs, H. T., Koenderink, J. J., & Kappers, A. M. L. (2005). The influence of illumination direction on the pictorial reliefs of Lambertian surfaces. *Perception*, 34(3), 275–287.
- Pentland, A. (1982). Finding the illuminant direction. *Journal of the Optical Society of America*, 72, 448–455.
- Pentland, A. (1989). Shape information from shading: A theory about human perception. *Spatial Vision*, 4, 165–182.
- Reichel, F. D., & Todd, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 653–664.
- Sun, P., & Schofield, A. J. (2012). Two operational modes in the perception of shape from shading revealed by the effects of edge information in slant settings. *Journal of Vision*, 12(1):12, 1–21, <http://www.journalofvision.org/content/12/1/12>, doi:10.1167/12.1.12. [PubMed] [Article]
- Todd, J. T., Egan, E. J. L., & Phillips, F. (2014). Is the perception of 3D shape from shading based on assumed reflectance and illumination? *I-Perception*, 5(6), 497–514.
- Todd, J. T., Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1996). Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, 22(3), 695–706.
- Todd, J. T., & Mingolla, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental*

- Psychology: Human Perception and Performance*, 9(4), 583–595.
- Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, 96(4), 643–657.
- Todorović, D. (2014). How shape from contours affects shape from shading. *Vision Research*, 103, 1–10.
- Wijntjes, M. W. A. (2012). Probing pictorial relief: From experimental design to surface reconstruction. *Behavior Research Methods*, 44(1), 135–143.
- Wijntjes, M. W. A., Doerschner, K., Kucukoglu, G., & Pont, S. C. (2012). Relative flattening between velvet and matte 3D shapes: Evidence for similar shape-from-shading computations. *Journal of Vision*, 12(1):2, 1–11, <http://www.journalofvision.org/content/12/1/2>, doi:10.1167/12.1.2. [PubMed] [Article]
- Zhang, R., Tsai, P. S., Creyer, J. E., & Shah, M. (1999). Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8), 690–706.