

Perceptual biases in the interpretation of 3D shape from shading

Baoxia Liu^a, James T. Todd^{b,*}

^a *University of California, Berkeley, USA*

^b *Department of Psychology, Ohio State University, 142 Townshend Hall, Columbus, OH 43210, USA*

Received 22 September 2003; received in revised form 10 March 2004

Abstract

Two experiments are reported in which observers judged the sign and magnitude of surface curvature from shaded images of an indoor scene. The depicted surfaces were illuminated by an area light in the ceiling or floor, and the illumination was attenuated with distance in a physically correct manner. The displays were presented both with and without cast shadows, specular highlights and surface inter-reflections in all possible combinations. The overall pattern of results revealed a strong perceptual bias to interpret the images as convex rather than concave, and a weaker bias to prefer illumination from above rather than from below, though there were large individual differences in the relative strengths of these biases. For displays that did not contain cast shadows or specular highlights, the accuracy of observers' judgments about the sign of surface curvature was no greater than chance, but performance was significantly improved when those aspects of normal shading were included in the rendering model. An analysis of the apparent depth magnitudes revealed that convex surfaces produce much greater perceived depth than concave surfaces with comparable relief.

© 2004 Elsevier Ltd. All rights reserved.

1. Introduction

One of the oldest and most famous illusions in visual perception is the apparent inversion of relief that can occur when a shaded image of a surface is viewed upside down (see Fig. 1). This phenomenon was first reported by Gmelin at the Royal Society of London in 1744, and it has been studied extensively since then by many other investigators (e.g., Benson & Yonas, 1973; Berbaum, Bever, & Chung, 1983, 1984; Brewster, 1826, 1832, 1847; Hagen, 1976; Hershberger, 1970; Hess, 1950; Ramachandran, 1988; von Fieandt, 1949). The classical explanation of why surfaces may appear inverted in depth was first proposed by the early American scientist David Rittenhouse (1786). He argued that observers are biased to interpret patterns of shading based on their prior knowledge about the direction of illumination. In the ecology of natural vision it is statistically more common for surfaces to be illuminated from above rather than from below. Thus, when presented with an otherwise ambiguous pattern of shading (see Belhum-

eur, Kriegman, & Yuille, 1999; Koenderink, van Doorn, Kappers, & Todd, 2001), observers will perceive the sign of surface relief that is consistent with an overhead illumination.

Although this explanation is often presented in textbooks on perception as if it were factually correct, many of the examples that have been used to demonstrate the inversion phenomenon are most likely based on other factors. During the past decade there has been a growing body of evidence to indicate that the perception of surface relief from shading may actually be influenced by several distinct biases, all of which are statistically well grounded in the ecology of the natural vision. For example, because of the existence of gravity, it is statistically more common to observe surfaces from above rather than from below, and this can have an important influence on the perception of 3D shape. Other things being equal, observers will perceive ambiguous images so that the overall pattern of surface depth increases with height in the image plane (Langer & Bülthoff, 2001; Mamassian & Landy, 1998; Reichel & Todd, 1990). This global orientation bias has been documented for surfaces depicted with shading, motion or texture (see Fig. 2), and we suspect it is responsible for most examples of the “crater illusion” like the one in Fig. 1. Another

* Corresponding author. Tel.: +1-614-292-8661; fax: +1-614-292-5601.

E-mail address: todd.44@osu.edu (J.T. Todd).




Fig. 1. Perceptual inversion of shaded surfaces. These two images are identical in all respects except that they are presented in opposite orientations. This change in orientation causes the apparent sign of relief to become inverted, so that the image on the left appears to have a small bump in its center, whereas the one on the right appears to have a small dimple.

important statistical regularity in the natural environment that can influence the perception of 3D shape involves the distribution of surface curvatures. Because all solid objects are globally convex, convex surface patches are statistically more common than concave surface patches, and observers can apparently exploit this regularity for interpreting the structure of ambiguous shaded images. That is to say, they are more likely to perceive surfaces as convex rather than concave (Hill & Bruce, 1994; Langer & Bühlhoff, 2001).

A particularly compelling demonstration of these biases has recently been reported by Langer and Bühlhoff (2001). They showed observers shaded images of globally convex or concave surfaces with a stucco-like texture that could face upward or downward and could be illuminated either from above or from below. On each trial a surface was presented together with a small probe point to mark one of its local regions, and observers were required to indicate whether the designated region appeared to be concave or convex. The results revealed that observers were biased to perceive the depicted surfaces as being globally convex, with a globally upward orientation and illuminated from above, and that all three of these biases had roughly the same strength. An especially interesting aspect of these results is that the overall accuracy of observers' judg-

ments was only 51%. This suggests that they were unable to make use of other available sources of information for determining the sign of curvature, such as shadows, occlusion contours or perspective, and that their judgments were determined entirely by perceptual biases.

In presenting their results, Langer and Bühlhoff raised an interesting caveat concerning the potential generality of these findings. A somewhat unusual aspect of their stimuli is that the surface undulations the observers were asked to judge had a much higher spatial frequency than has been used by other researchers for the study of 3D shape from shading, and this may have affected the relative detectability of some possible sources of information about the local pattern of relief. For example, when the scale of surface structure becomes sufficiently small, it may be difficult to reliably distinguish cast and attached shadows, or to identify the attached sides of smooth occlusion contours.

In an effort to further address this issue, the research described in the present article was designed to examine the perception of surface curvature from photo-realistic shading patterns of a large scale visual environment. The goals of this research were twofold: First, to assess how the perceived sign and magnitude of surface curvature are influenced by the presence or absence of potential information from cast shadows, specular highlights, indirect illumination and smooth occlusion contours; and second, to evaluate the relative strengths of the overhead illumination and global convexity biases with varying amounts of visual information that could potentially specify the correct sign of curvature.

2. Experiment 1

2.1. Methods

2.1.1. Apparatus

The experiment was conducted using a Dell Precision 420 PC with a GeForce3 graphics card and a 53.34 cm




Fig. 2. Perceptual inversion of textured surfaces. These two images are identical in all respects except that they are presented in opposite orientations. Although there is no smooth shading in these images, the change in orientation causes the apparent sign of relief to become inverted, in exactly the same way as in Fig. 1. This is most likely due to a perceptual bias to interpret the overall surface slant so that depth increases with height in the image plane (see Reichel & Todd, 1990).

monitor that was gamma corrected to linearize the display intensities. Stimuli were presented within a 37.5×30.0 cm rectangular region of the display screen with a spatial resolution of 1280×1024 pixels at refresh frequency of 85 Hz. The displays were viewed monocularly from a distance of 114 cm, such that the display region subtended $18.75^\circ \times 15.0^\circ$ of visual angle.

2.1.2. Stimuli

The stimuli depicted the inside of a $304.8 \times 304.8 \times 304.8$ cm room that was created using 3D StudioMax 4.2. The center of one wall contained an ellipsoidal concavity or convexity with a height and width of 101.6 cm and a depth that varied across trials with possible values of 35.56, 50.8 and 66.04 cm. The viewing position was located in the center of the opposite wall. This scene was illuminated by a 50.8×50.8 cm square area light that could be located in the center of the ceiling or the floor.

The images were rendered in Lightscape 3.2 using both radiosity and ray-tracing algorithms. Lightscape uses a surface reflectance model developed originally by Torrance and Sparrow (1967) that provides a close match to photometric measurements of how light reflects from real physical surfaces. The material properties in the present experiment were controlled by three parameters: reflectance, shininess and the index of refraction. Reflectance refers to the proportion of incident light that is reflected. Shininess refers to the smoothness

of the surface at a microscopic scale—i.e. it is the proportion of surface micro facets that face in the direction of the average surface normal. The index of refraction defines the amount of light that enters the surface material as opposed to being reflected off its surface. The higher the value of the index of refraction, the less light will be transmitted to the interior of the material, and a value of one means that all of the light will be transmitted. Most natural materials have an index of refraction between 1.0 and 1.5.

For the scenes used in the present experiments, the ellipsoidal region and its background plane could be either matte or glossy, and all of the other surfaces in the room were matte. The matte surfaces had a reflectance of 0.7, a shininess of 0.0 and a refractive index of 1.0. The glossy surfaces had a reflectance of 0.7, a shininess of 0.5 and a refractive index of 1.25. These surfaces could be presented either with or without visible cast shadows or indirect illumination from surface inter-reflections. For those displays with indirect illumination, the radiosity algorithm was allowed to iterate until 99% of the light had been dissipated. In order to approximately equate the luminance ranges among the different conditions the simulated light source intensity was adjusted for different parameter settings over a possible range from 3000 to 6000 cd. Fig. 3 shows some example images of convex and concave surfaces both with and without cast shadows, specular highlights and indirect illumination from surface inter-reflections.




Fig. 3. Some example images from Experiment 1. The two upper panels depict convex surfaces, and the two lower panels depict concave surfaces. The images on the right contain cast shadows, specular highlights and surface inter-reflections, whereas these properties have been removed from the images on the left.

2.1.3. Procedure

On each trial a 512×512 pixel image of an ellipsoidal concavity or convexity was presented on the left side of the computer monitor. The right side of the monitor contained an elliptical curve that observers could adjust with a hand held mouse to indicate the apparent sign and magnitude of surface curvature in each condition (see Fig. 4). When satisfied with their settings, observers pressed the space bar to initiate a new trial. They received no feedback about their performance until after the experiment was completed.

2.1.4. Subjects

Seven observers participated in the experiment, including the two authors and five others who were naïve about the issues being investigated. They all had normal or corrected-to-normal visual acuity.

2.1.5. Design

To summarize the overall experimental design, there were 96 possible conditions: 2 signs of curvature (concave or convex) \times 3 depth magnitudes (35.56, 50.8 and 66.04 cm) \times 2 levels of shininess (matte or glossy) \times 2 illumination models (with or without surface inter-reflections) \times 2 illumination directions (from above or below) \times 2 types of shadowing (with or without cast shadows). Within a given experimental session, each of




Fig. 4. The adjustment task for Experiment 1. In order to indicate the perceived sign and magnitude of curvature, observers adjusted a curve presented adjacent to each shaded image so that it matched the apparent depth profile of the depicted surface.

the 96 possible shaded images was presented in a random sequence. Each subject participated in five separate sessions.

2.2. Results

Before describing the results of this study, it is useful to point out some potential sources of information in Fig. 3 that could be used to determine the direction of illumination and the sign of surface curvature. The direction of illumination in these images is optically specified by overall direction of the light attenuation gradient, and, for glossy surfaces, by the positions of the specular highlights. When the scene is illuminated from above, for example, the local maxima of the diffuse and specular components of shading are both located on the upper portion of the background wall, whereas they are both located on the lower portion of the background wall when the scene is illuminated from below. There is other information to specify the sign of curvature of the ellipsoidal regions. For convex surfaces, the cast shadows are located on the background wall, whereas for concave surfaces, they are contained within the circular boundary of the ellipsoidal region. For images of convex surfaces, the luminance gradient of the ellipsoidal region is in the same direction as on the background wall, whereas it is in the opposite direction for images of concave surfaces. Similarly, for the images of convex glossy surfaces, the highlight on the ellipsoidal region is on the side nearest the highlight on the background wall, whereas it is on the opposite side for images of concave surfaces.

Were the observers in this study able to make use of the available information for accurately specifying the sign of surface curvature, or, did they rely instead on perceptual biases as in the earlier study by Langer and Bühlhoff (2001)? When averaged over all observers and conditions the percentage of correct sign of curvature judgments was only 56%—just barely above chance. Thus, it appears that the available information had only a minimal impact on whether the surfaces were perceived as concave or convex.

Fig. 5 provides a more detailed breakdown of how various stimulus factors influenced the apparent sign of curvature in these displays. It is interesting to note that the presence or absence of surface inter-reflections or specular highlights had no effects whatsoever on the perceived sign of curvature. Performance was improved somewhat by the addition of cast shadows, but this only increased the accuracy of observers' judgments to a level of 63%. It is clear from these data that the judged pattern of relief was primarily determined by the observers' perceptual bias for global convexity, because 80% of the surfaces were identified as convex. There was also a much smaller bias for overhead illumination, which accounted for approximately 60% of the responses.




Fig. 5. The percentage of correct sign of curvature judgments in Experiment 1 for each sign of curvature and direction of illumination, and for the presence or absence of cast shadows, specular highlights or surface inter-reflections. The error bars show 95% confidence intervals.




Fig. 6. The percentage of responses for each individual observer in Experiment 1 that were consistent with a convexity bias or a bias to perceive illumination from above.

Additional analyses revealed that there were significant individual differences among observers in the relative strengths of these biases. Fig. 6 shows the percentage of the total responses for each observer that was consistent with a bias for overhead illumination or global convexity. Note that four of the seven observers had a strong bias for global convexity without any preference for the direction of illumination. One observer produced the opposite pattern of performance, with a strong bias for overhead illumination and no preference for sign of curvature. One observer exhibited both biases, and the remaining one exhibited neither.

An important component in the design of this study is that we also measured the magnitude of perceived depth

in addition to the apparent sign of curvature. In analyzing these data, the observers' depth judgments were represented as unsigned quantities irrespective of the judged sign of curvature. Fig. 7 shows the average absolute value of judged depth as a function of the simulated depth for each combination of shadowing, glossiness and sign of relief. An analysis of variance of these data revealed that the apparent depth magnitudes of the convex surfaces were significantly increased by the magnitude of simulated depth, $F(2, 12) = 54.06$, $p < 0.001$, the presence of cast shadows, $F(1, 6) = 14.65$, $p < 0.01$, and the presence of specular highlights, $F(1, 6) = 50.259$, $p < 0.001$ (see also Todd & Mingolla, 1983; Todd, Norman, Koenderink, & Kappers, 1997). A surprising result we had not anticipated is that the apparent depths of the concave surfaces appeared much smaller than the convex surfaces, $F(1, 6) = 101.02$, $p < 0.001$, and they were statistically independent of the magnitude of simulated depth. The judged depths of the concave surfaces were significantly increased by the presence of specular highlights, $F(1, 6) = 17.76$, $p < 0.001$, and they were slightly decreased by the presence of cast shadows, $F(1, 6) = 12.60$, $p < 0.05$.

3. Experiment 2

In evaluating the results of Experiment 1, it is interesting to note that there is one potential factor that may have artificially inflated the percentage of surface convexity judgments. Previous research has shown that smooth occlusion contours provide a powerful source of information about the sign of surface curvature in their immediate local neighborhoods. It can be shown mathematically that the sign of surface curvature in a direction perpendicular to an attached occlusion contour must always be convex (Koenderink, 1984;




Fig. 7. The average absolute value of judged depth in Experiment 1 as a function of the simulated depth for each combination of shadowing, glossiness and sign of relief. The standard error in all of the depicted conditions was approximately 1.5 cm.

Koenderink & van Doorn, 1982), and there have been several empirical studies to show that this information can influence observers' perceptions of 3D shape from shading (Howard, 1983; Reichel & Todd, 1990; Todd & Reichel, 1989). The reason this may be relevant to the results of Experiment 1 is that it is possible that the sharp edges of the concave surface regions may have been perceptually misinterpreted as smooth occlusion contours. If so, then those contours may have provided misleading information that the depicted surface regions were convex. In Experiment 2 we attempted to remove

this possible misinterpretation by presenting convex and concave surface regions with smoothed edges.

3.1. Methods

The apparatus, procedure, design and observers were the same as in Experiment 1. The only difference was that the convex and concave surface regions had a bell-shaped rather than an ellipsoidal structure (see Fig. 8), and the schematic depth profile that the subjects ad-




Fig. 8. Some example images from Experiment 2. The two upper panels depict convex surfaces, and the two lower panels depict concave surfaces. The images on the right contain cast shadows, specular highlights and surface inter-reflections, whereas these properties have been removed from the images on the left.




Fig. 9. The adjustment task for Experiment 2. In order to indicate the perceived sign and magnitude of curvature, observers adjusted a curve presented adjacent to each shaded image so that it matched the apparent depth profile of the depicted surface.

justed to indicate their perceptions was modified accordingly (see Fig. 9).

4. Results

Fig. 10 provides a breakdown of how various stimulus factors influenced the apparent sign of curvature in these displays, and Fig. 11 shows the relative strengths




Fig. 11. The percentage of responses for each individual observer in Experiment 2 that were consistent with a convexity bias or a bias to perceive illumination from above. The error bars show 95% confidence intervals.

of the overhead illumination and convexity biases for each individual observer. In general, the results were quite similar to those obtained in Experiment 1, though there were two interesting differences that deserve to be noted. First, whereas the highlights in Experiment 1 had no detectable influence on the judged sign of curvature, their presence in these displays produced a 10% increase in the accuracy of observers' judgments. This is most likely due to the placement of the highlights along the smoothly curved boundaries of these surfaces, which did not occur on the sharp edges of the ellipsoidal surface patches in Experiment 1 (see Figs. 3 and 8).




Fig. 10. The percentage of correct sign of curvature judgments in Experiment 2 for each sign of curvature and direction of illumination, and for the presence or absence of cast shadows, specular highlights or surface inter-reflections. The error bars show 95% confidence intervals.




Fig. 12. The average absolute value of judged depth in Experiment 2 as a function of the simulated depth for each combination of shadowing, glossiness and sign of relief. The standard error in all of the depicted conditions was approximately 1.4 cm.

The second interesting difference between Experiments 1 and 2 involves the relative strength of observers' biases to perceive surfaces as convex rather than concave. Whereas 80% of the surfaces in Experiment 1 were judged to be convex, only 70% of the surfaces appeared convex in Experiment 2. Note in Figs. 5 and 10 that this reduction in apparent convexity only occurred for the concave surfaces. This finding is consistent with the hypothesis that the sharp edges of the concave surface regions in Experiment 1 may have been perceptually misinterpreted as smooth occlusion contours, which would have provided misleading information that these surfaces were convex. When the sharp edges were removed in the present study, the observers' accuracy at judging concave surfaces increased from 26% to 43%. It is important to keep in mind, however, that this increased accuracy was limited to those displays that contained cast shadows or specular highlights. When the displays included both of those features, the overall level of performance in judging the sign of curvature was 71%. When neither of those features was present, the accuracy of the observers' judgments dropped to 49%, which is not significantly different from chance.

Fig. 12 shows the average absolute value of judged depth as a function of the simulated depth for each combination of shadowing, glossiness and sign of relief. These results are similar in most respects to those obtained in Experiment 1, except that the manipulation of cast shadows had no significant effect on performance. As in the previous study, the convex surfaces appeared to have much greater depth than the concave surfaces, $F(1, 6) = 39.94$, $p < 0.001$. For convex surfaces, the magnitude of perceived depth was significantly increased by the presence of specular highlights, $F(1, 6) = 25.37$, $p < 0.01$, and the magnitude of simulated depth, $F(2, 12) = 59.20$, $p < 0.001$. The apparent depth magnitudes of concave surfaces were also increased by the presence of specular highlights, $F(1, 6) = 24.52$,

$p < 0.01$, but, as in Experiment 1, the perceived depths of those displays were unaffected by the magnitude of simulated depth.

5. General discussion

There are many different factors that can influence patterns of shading in natural vision. For example, within indoor environments like the ones depicted in the present experiments, the intensity of direct illumination at each point decreases as the inverse square of its distance from the light source, which produces visible light attenuation gradients. This same inverse square law is also applicable in outdoor scenes, but because of the enormous distance to the sun, the attenuation gradients are negligible at scales that are relevant to human perception. Variations in viewing conditions can also affect the nature of cast shadows—ranging from soft shadows with relatively large penumbras in many indoor environments, to hard shadows with no visible penumbras in outdoor scenes on bright sunny days. Because of these contextual variations in patterns of image shading it is quite possible that potential sources of information that are available in one context may not be present in another.

In a recent experiment on the perception of surface curvature from shading Langer and Bühlhoff (2001) examined the effects of three different types of perceptual bias: A bias to perceive surfaces as globally convex, a bias to perceive surfaces as illuminated from above, and a bias to perceive surfaces as viewed from above. The displays used in their study depicted surfaces with a stucco-like texture that contained many different regions of concavity and convexity, and they were illuminated by a pattern of collimated light as would typically be encountered in an outdoor scene on a sunny day. Although the simulation they employed was more eco-

logically valid than those used in many previous studies of the perception of shape from shading, their results were quite remarkable in that the accuracy of observers' judgments about the sign of curvature in different local regions was no greater than chance. This finding suggests that observers were unable to make use of the available sources of information for determining the sign of curvature, and that they relied almost entirely on perceptual biases in making their judgments.

Are human observers truly insensitive to potential information in shading about the sign of surface curvature, or is there something about the context used by Langer and Bühlhoff (2001) that made this information difficult to detect? The research described in the present article was designed in part to address this issue. It is important to note that the surface geometry used in our study was much simpler than the one used by Langer and Bühlhoff. Each display contained a single concave or convex surface patch that was imbedded within the back wall of a $304.8 \times 304.8 \times 304.8$ cm room. The pattern of illumination was more complex, however. The surfaces were illuminated by a 50.8×50.8 cm square area light that was located in the center of the ceiling or the floor, which produced patterns of shading with visible light attenuation gradients, cast shadows, specular highlights and global illumination from surface inter-reflections. For the displays that did not contain cast shadows or specular highlights, our results were identical to those reported by Langer and Bühlhoff. That is to say, the ability of observers' to determine the correct the sign of curvature was no greater than chance. However, the level of performance was significantly improved when cast shadows and specular highlights were present, thus suggesting that these visual features can provide perceptually useful information about the sign of surface curvature under at least some viewing conditions.

Although observers' performance was significantly greater than chance for displays that contained cast shadows and specular highlights, the overall pattern of results is, nonetheless, largely consistent with the conclusions of Langer and Bühlhoff. Even in the most optimal conditions, the highest level of accuracy achieved was only 71%, and it is difficult to imagine another possible environment that would provide more information than our displays without introducing other modalities such as motion or stereo (e.g., Hill & Bruce, 1993, 1994). It appears to be the case, therefore, that the apparent sign of curvature from shading is heavily influenced by perceptual biases even under ideal conditions for detecting the available information.

In evaluating the relative strengths of different perceptual biases, Langer and Bühlhoff found that the bias for global convexity is roughly 30% larger than the bias for overhead illumination. In our experiments this difference was much larger, though there are two important factors that complicate the analysis of this issue:

One is the fact that that there were large individual differences among our subjects (see Figs. 6 and 11), which makes it dangerous to generalize to the entire population. A second complicating factor is that the sharp boundaries of the concave regions in Experiment 1 may have been mistaken for smooth occlusion contours, which would have provided misleading information that the depicted surface regions were convex (see Howard, 1983; Koenderink, 1984; Koenderink & van Doorn, 1982; Reichel & Todd, 1990; Todd & Reichel, 1989). When these sharp edges were removed in Experiment 2 the strength of the convexity bias was significantly reduced.

Given the large body of literature on overhead illumination biases in the perception of 3D shape from shading, it is interesting to speculate why many of the observers in the present studies showed no bias at all with respect to the direction of illumination. Based on our own perceptual experiences with shaded images, we suspect it is the case that the overhead illumination bias is a relatively weak phenomenon that is easily overridden by other perceptual biases or available sources of information about the sign of surface relief. For example, one of the most common demonstrations that is purported to demonstrate such a bias is the crater illusion as exemplified in Fig. 1, but this effect is most likely due to a bias to perceptually interpret surfaces as viewed from above rather than viewed from below (Langer & Bühlhoff, 2001; Mamassian & Landy, 1998, 2001; Reichel & Todd, 1990). The most compelling evidence for an overhead illumination bias in human perception comes from experiments with fronto-parallel surfaces containing equal amounts of concavity and convexity so that other perceptual biases are neutralized (eg., Benson & Yonas, 1973; Berbaum et al., 1983, 1984; von Fieandt, 1949; Kleffner & Ramachandran, 1992; Mamassian & Goutcher, 2001; Ramachandran, 1988). Some typical examples that are representative of these studies are shown in the upper panels of Fig. 13, which show a pattern of bumps and dents under two different lighting conditions. For most observers, the apparent sign of relief is reversed if these images are viewed upside down or if their gray scale values are inverted. It is also important to note, however, that the conditions for creating this effect are highly constrained. The apparent depth inversion is most easily achieved under natural lighting conditions for surfaces whose concavities and convexities are sufficiently shallow that they do not produce cast shadows. The lower two panels of Fig. 13 show a surface with hemispherical bumps and dents that does not conform with this restriction. Note how the presence of visible cast shadows makes the surface more resistant to perceptual reversal when the image is viewed upside down.

The ability of human observers in the present experiments to use cast shadows and specular highlights as




Fig. 13. Shaded images of surface concavities and convexities with varying magnitudes of relief and varying patterns of illumination. The depth profiles of the depicted surfaces are shown in the right column. The direction of illumination in each case is at a 45° angle relative to the image plane. The surface depicted in the upper two panels has shallow concavities and convexities that do not produce cast shadows. Many observers report that the apparent sign of relief is reversed if these images are viewed upside down or if their gray scale values are inverted. The lower two panels show a surface with a greater magnitude of relief that does produce cast shadows. Note in this case that the apparent 3D structure is more resistant to perceptual reversal when the images are viewed upside down.

sources of information for the interpretation of shaded images is consistent with several other studies that have examined the perceptual effects of these optical phenomena in a variety of different contexts (e.g., Berbaum et al., 1984; Blake & Bülthoff, 1991; Bülthoff & Mallot, 1988; Cavanagh & Leclerc, 1989; Erens, Kappers, & Koenderink, 1993; Knill & Kersten, 1991; Madison, Thompson, Kersten, Shirley, & Smits, 2001; Mamassian, Knill, & Kersten, 1998; Norman, Todd, & Orban, in press; Todd & Mingolla, 1983; Todd et al., 1997; Todd, Norman, & Mingolla, 2004; Wanger, Ferwerda, & Greenberg, 1992). Although there was no evidence in the present study that indirect light from surface inter-reflections provides useful information for the perception of surface curvature, other research has shown that indirect illumination can facilitate judgments of lightness (Gilchrist & Jacobsen, 1984), or determining whether a surface is in contact with the ground (Madison et al., 2001). It is interesting to note that there are no current theoretical models that can adequately account for any of these findings. Existing computational algorithms for determining 3D shape from shading are typically based on strong a priori assumptions that all surfaces in a scene have matte reflectance functions and that the pattern of illumination is spatially homogeneous (e.g., see Horn & Brooks, 1989). It is clear from the empirical evidence,

however, that human perception of 3D shape from shading is considerably more robust.

Perhaps the most surprising aspect of the present experiments is the large asymmetry in the relative perceived depths of the convex and concave surface patches. All of the observers reported that the concave surfaces appeared to have little depth at all, and their judgments for those surfaces were completely independent of the magnitude of simulated depth. It is not immediately obvious why the perception of shape from shading should break down for concave surface patches. One possibility is that ambient light becomes trapped within concave regions, which reduces the overall luminance contrast, but the same effect was also obtained when the ambient light was turned off. Another possibility is that the relevant information for specifying a concavity competes with the perceptual bias for convexity, which results in a flattened compromise interpretation. Additional research will obviously be required in order to provide a more complete explanation of this phenomenon.

Acknowledgements

This research was supported by grants from NIH (R01-Ey12432) and NSF (BCS-0079277).

References

- Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision*, *35*, 33–44.
- Benson, C. W., & Yonas, A. (1973). Development of sensitivity to static pictorial depth information. *Perception & Psychophysics*, *13*, 361–366.
- Berbaum, K., Bever, T., & Chung, C. S. (1983). Light source position in the perception of object shape. *Perception*, *12*, 411–416.
- Berbaum, K., Bever, T., & Chung, C. S. (1984). Extending the perception of shape from known to unknown shading. *Journal of the Optical Society of America A*, *4*, 1155–1167.
- Blake, A., & Bülthoff, H. H. (1991). Shape from specularities: computation and psychophysics. *Philosophical Transactions of the Royal Society of London B*, *331*, 237–252.
- Brewster, D. (1826). On the optical illusion of the conversion of cameos into intaglios, and of intaglios into cameos, with an account of other analogous phenomena. *Edinburgh Journal of Science*, *4*, 99–108.
- Brewster, D. (1832). *Letters on natural magic*. London: John Murray.
- Brewster, D. (1847). On the conversion of relief by inverted vision. *Transactions of the Royal Society of Edinburgh*, *15*, 657–662.
- Bülthoff, H. H., & Mallot, H. A. (1988). Integration of depth modules: stereo and shading. *Journal of the Optical Society of America A*, *5*, 1749–1758.
- Cavanagh, P., & Leclerc, Y. G. (1989). Shape from shadows. *Journal of Experimental Psychology Human Perception and Performance*, *15*, 3–27.
- Erens, R. G. F., Kappers, A. M. L., & Koenderink, J. J. (1993). Perception of local shape from shading. *Perception & Psychophysics*, *54*, 145–156.
- Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception*, *13*, 5–19.
- Hagen, M. A. (1976). The development of sensitivity to cast and attached shadows in pictures as information for the direction of the source of illumination. *Perception & Psychophysics*, *20*, 25–28.
- 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.
- Hill, H., & Bruce, V. (1993). Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion. *Perception*, *22*, 887–897.
- Hill, H., & Bruce, V. (1994). A comparison between the hollow-face and "hollow-potato" illusions. *Perception*, *23*, 1335–1337.
- Horn, B. K. P., & Brooks, M. J. (1989). *Shape from shading*. Cambridge, MA: MIT Press.
- Howard, I. P. (1983). Occluding edges in apparent reversal of convexity and concavity. *Perception*, *12*, 85–86.
- Kleffner, D. A., & Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics*, *52*, 18–36.
- Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature*, *351*, 228–230.
- Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, *13*, 321–330.
- Koenderink, J. J., & van Doorn, A. J. (1982). van The shape of smooth objects and the way contours end. *Perception*, *11*, 129–137.
- Koenderink, J. J., van Doorn, A. J., Kappers, A. M., & Todd, J. T. (2001). Ambiguity and the "mental eye" in pictorial relief. *Perception*, *30*, 431–448.
- Langer, M. S., & Bülthoff, H. H. (2001). A prior for local convexity in local shape from shading. *Perception*, *30*, 403–410.
- Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001). Use of interreflection and shadow for surface contact. *Perception & Psychophysics*, *63*, 187–194.
- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*, *2*, 288–295.
- Mamassian, P., & Goutcher, R. (2001). Prior knowledge on the illumination position. *Cognition*, *81*, B1–B9.
- Mamassian, P., & Landy, M. S. (1998). Observer biases in the interpretation of line drawings. *Vision Research*, *38*, 2817–2832.
- Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, *41*, 2653–2668.
- Norman, J. F., Todd, J. T., & Orban, G. A. (in press). Perception of 3D shape from specular highlights and deformations of shading. *Psychological Science*, *15*.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, *331*, 163–166.
- 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.
- Rittenhouse, D. (1786). Explanation of an optical deception. *Transactions of the American Philosophical Society*, *2*, 37–42.
- Todd, J. T., & Mingolla, E. (1983). The perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 583–595.
- Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. L. (1997). Effects of texture, illumination and surface reflectance on stereoscopic shape perception. *Perception*, *26*, 806–822.
- Todd, J. T., Norman, J. F., & Mingolla, E. (2004). Lightness constancy in the presence of specular highlights. *Psychological Science*, *15*, 33–39.
- Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, *96*, 643–657.
- Torrance, K., & Sparrow, E. (1967). Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, *57*, 1105–1114.
- von Fieandt, K. (1949). The phenomenological problem of light and shadow. *Acta Psychologica*, *6*, 337–357.
- Wanger, L. R., Ferwerda, J. A., & Greenberg, D. P. (1992). Perceiving spatial relationships in computer-generated images. *IEEE Computer Graphics and Applications*, *12*, 44–55.