

The Assumed Light Direction for Perceiving Shape from Shading

James P. O’Shea*
Vision Science
University of California, Berkeley

Martin S. Banks†
Vision Science
University of California, Berkeley

Maneesh Agrawala‡
Computer Science
University of California, Berkeley




Figure 1: The shading of an object’s surface depends on the light source direction. The light is directed from the viewpoint in the left image, from 22 deg above the viewpoint in the middle image, and from 44 deg above the viewpoint in the right image. The object is positioned identically in each of the three views. In this paper, we present an experiment designed to test how shape perception is affected by changing the angle of the light direction. We found the lighting used in the center image led to the most accurate estimations of 3D shape.

Abstract

Recovering 3D shape from shading is an ill-posed problem that the visual system can solve only by making use of additional information such as the position of the light source. Previous research has shown that people tend to assume light is above and slightly to the left of the object [Sun and Perona 1998]. We present a study to investigate whether the visual system also assumes the angle between the light direction and the viewing direction. We conducted a shape perception experiment in which subjects estimated surface orientation on smooth, virtual 3D shapes displayed monocularly using local Lambertian shading without cast shadows. We varied the angle between the viewing direction and the light direction within a range ± 66 deg (above/below), and subjects indicated local surface orientation by rotating a gauge figure to appear normal to the surface [Koenderink et al. 1992]. Observer settings were more accurate and precise when the light was positioned above rather than below the viewpoint. Additionally, errors were minimized when the angle between the light direction and the viewing direction was 20-30 deg. Measurements of surface slant and tilt error support this result. These findings confirm the light-from-above prior and provide evidence that the angle between the viewing direction and the light direction is assumed to be 20-30 deg above the viewpoint.

Keywords: Shape-from-shading, light-from-above, shape perception, Lambertian shading.

*e-mail: joshea@cs.berkeley.edu

†e-mail: martybanks@berkeley.edu

‡e-mail: maneesh@cs.berkeley.edu

1 Introduction

Shading is a fundamental pictorial depth cue. In computer graphics, shading is typically employed to convey 3D shape on a flat display screen. This type of visualization is completely dependent on the visual system’s ability to perceive shape from a shaded 2D image. In order to convey 3D information effectively, it is therefore important to understand how shading influences shape perception.

Shading conveys depth by varying the levels of brightness in an image to indicate the extent to which a surface is oriented towards a light source. In computer graphics, there are many methods to produce this effect with varying degrees of realism and complexity [Torrance and Sparrow 1967; Phong 1975; Cook and Torrance 1982; Ward 1992; Oren and Nayar 1994]. Consider the image of a surface with an albedo ρ . In a simple diffuse shading model, the intensity I of any particular point in the image is proportional to the cosine of the angle between the surface normal \hat{n} and the vector pointing to the light \hat{l} as follows:

$$I = \rho(\max(\hat{n} \cdot \hat{l}, 0)) \quad (1)$$

Computing the 3D shape from this image requires solving for the surface normals \hat{n} . There is no unique local solution for this formulation of the shape-from-shading problem because we have one

Copyright © 2008 by the Association for Computing Machinery, Inc.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.

APGV 2008, Los Angeles, California, August 9–10, 2008.

© 2008 ACM 978-1-59593-981-4/08/0008 \$5.00

equation with more than one unknown variable [Oliensis 1991; Prados and Faugeras 2005]. Additionally, Belhumeur and colleagues showed that the image of a shaded Lambertian surface is consistent with a continuous set of related surfaces [Belhumeur et al. 1997]. Due to this bas-relief ambiguity, the 3D shape of a Lambertian surface cannot be uniquely solved with a single viewpoint and an unknown light source.

One way the visual system can help resolve these ambiguities is therefore to assume the direction of light. We define the light direction according to a coordinate system in which the eye is positioned along the positive z-axis, the x-axis is horizontal (positive to the right), and the y-axis is vertical (positive up). We can therefore describe the direction of the light as a vector in 3-space (Figure 2). Assuming a light source at infinity, we can ignore the magnitude of this vector and specify the light direction using two parameters in a manner analogous to surface slant and tilt [Stevens 1983]. Surface slant is the angle between the viewing direction and the surface normal. Surface tilt is the angular difference between the positive x-axis and the surface normal after it is projected onto the xy-plane. We can similarly define a light direction according to its slant and tilt relative to the viewing direction. Lighting slant, in this case, is the angle between the viewing direction (z-axis) and the light direction vector.

Previous research has shown that the visual system assumes light is coming from above and slightly to the left of a shaded object [Sun and Perona 1998; Mamassian and Goutcher 2001]. Several studies have investigated how lighting and shading affect shape perception [Mamassian and Goutcher 2001; Mamassian and Kersten 1995; Koenderink et al. 1996; Langer and Bulthoff 2001], yet none have specifically examined variations in the slant of light.

We conducted a study to determine if changes in the slant of light affect 3D shape perception when shading is the only information available. We instructed observers to estimate the local surface orientation at specified sites on a virtual 3D object by orienting a gauge figure to appear normal to the surface [Koenderink et al. 1992]. By varying the light slant throughout the experiment, we were able to analyze subject settings as a function of illumination direction to determine the direction at which performance was optimal. If the visual system assumes a particular angle between the light direction and the view direction, we expect subject settings to be most accurate when the actual lighting is consistent with this assumption.

We rendered the surfaces using a local Lambertian shading model. We did not render specular reflections or cast shadows. We used this model because of its simplicity and its common usage in computer graphics applications. Due to its lack of photorealism however, the results of our study do not necessarily apply to viewing objects in the real world. Our findings do have a direct application in computer graphics. Knowledge of the assumed light position is relevant to recent work on automatic lighting design [Shaded and Lischinski 2001; Gumhold 2002; Lee and Hao 2006], as well as non-photorealistic rendering techniques designed to affect the perception of 3D shape [DeCarlo et al. 2003; Rusinkiewicz et al. 2006]. If such a system positions the light away from the assumed direction, a viewer may be less likely to accurately perceive the shape of the object.

The main contributions of this paper include the following:

- We measure shape perception as a function of the slant of light in order to directly evaluate whether the visual system assumes a particular slant angle when presented with local Lambertian shading without cast shadows. We find that observers assume the light is slanted 20-30 deg above the viewpoint.




Figure 2: *Lighting directions tested in the experiment. We varied the angle between the light source and the view vector within a range +/- 66 deg above and below the viewpoint in 11-deg increments. Previous research has shown that people assume the light is above and to the left an average of 12 deg from vertical. Thus, we constrain our light directions to a plane that has been rotated about the z-axis 12 deg counter-clockwise. The images to the left of the light source directions illustrate how an object would be shaded according to these light directions. The asterisks denote the light directions which led to the most accurate estimations of 3D shape in our study.*

- We isolate shading from other cues to shape, such as silhouettes and familiarity, in order to examine how shading alone affects the perception of surface orientation.

2 Related Work

To recover 3D shape from the image of a shaded surface, the visual system often makes use of information other than shading. Mamassian and Kersten demonstrated that the occluding contour of an object can be a significant cue to surface orientation [Mamassian and Kersten 1995]. In their experiment, subjects were instructed to orient a gauge figure to estimate local surface orientation for an object that is shaded under various lighting conditions. This work provides evidence that shading does affect surface perception, but it is important to note that their observer responses were not significantly different when the shading information was completely omitted and only the silhouette remained. The authors conclude that observers use the occluding contour of the object rather than shading to estimate surface orientation.

This finding may be due to Mamassian and Kersten's use of a simple object for probing local surface perception. Its surface was relatively uniform, and the silhouette provided adequate information for estimating surface orientation. By including the silhouette, this study failed to investigate the effect of shading alone. Given the limited number of lighting conditions, their experiment was also not specifically designed to examine how varying the angle between the view vector and the light direction would affect surface perception.

The familiarity of a shape may also provide additional information for estimating local surface orientation. Koenderink et al. [1996] investigated how shading affects surface perception using a gauge-figure technique, yet their use of the human form as a stimulus image may have confounded the results. In the same way that the occluding contour of an object provides additional shape information beyond the shading, an observer’s familiarity with an object may influence the results in this type of experiment. Because Koenderink’s study was not specifically designed to address this question, it is unclear how it may have affected the outcome.

There is also research suggesting that the global convexity of a shape may influence the perception of local surface orientation for shaded objects. Langer and Bühlhoff [2001] showed that performance in a surface discrimination task was better for surfaces of globally convex objects compared to surfaces from globally concave objects. Their surfaces were rendered using only three different lighting directions.

The light-from-above prior is an important assumption a viewer can make to disambiguate the shape of a shaded surface. Sun and Perona experimentally investigated this prior in a study in which subjects were asked to identify if shaded ellipses appeared convex or concave [Sun and Perona 1998]. By rotating the shading gradient of elliptical discs about the visual axis, they effectively rotated the location of the light source as well. The results were consistent with an above-left prior. The study was not designed to specifically probe the angle between the light direction and the viewing direction.

This light-from-above prior was confirmed with line-drawing stimuli meant to represent the embossed surface of a fronto-parallel object [Mamassian and Goutcher 2001]. In these stimuli, dark lines indicated a part of the surface facing down and white lines indicated the highlights of a surface facing up. Given the simplicity of these stimuli, the lighting could only be perceived to be in one of two positions: above or below the pattern of lines. The authors of this study did not specifically address whether task performance was affected by the slant of light.

More recently, Caniard and Fleming [2007] addressed how shape estimation can be affected by changes in illumination conditions. They instructed subjects to perform both a surface-matching task, as well as local surface estimation task, while varying the location of the light source. Their results show an effect of light position, but they did not specifically vary the slant of the light in a systematic way.

These studies examined how lighting and shading affect shape perception. They did not specifically vary the angle between the light direction and the view direction, or they did not eliminate other pictorial cues which an observer could use to disambiguate the shape of a shaded surface. The experiment we describe in this paper addresses both of these concerns.

3 Methods

We used a gauge-figure task [Koenderink et al. 1992] to test surface perception while varying the angle between the light direction and the viewing direction. We presented images of shaded irregular surfaces on a computer display and subjects were instructed to orient a gauge figure to appear normal to the surface of the displayed object.

3.1 Subjects

We collected data from three subjects familiar with psychophysical tasks. None had seen the stimuli before or were aware of the specific hypothesis under investigation. All subjects had normal visual




Figure 3: Example screenshot from the experiment. In this figure, we show a synthetic aperture to approximately demonstrate how the stimuli appeared to the subject. In the experiment, we positioned a physical aperture between the viewer and the display screen to guarantee that the subjects would not mistake the edges of the aperture for the circular silhouette of the object.

acuity. If they required an optical correction, they wore it during testing.

3.2 Setup

The experiment was conducted using a desktop computer running Windows 2000 on an Intel processor. The 19-inch CRT display was gamma-corrected to linearize the luminance function for each color channel. We tested each subject monocularly after positioning them 45cm from the screen with their viewing eye directly in front of the center of the display. The subject’s head position was maintained throughout the experiment using a custom-fitted bite bar. The experiment was self-paced. Subjects responded using a mouse. The CRT was viewed through a physical aperture to prevent the observer from seeing and thereby making use of the occluding contour to judge surface orientation. The experiment was conducted in a dark environment.

3.3 Stimuli

We generated smooth irregular 3D models for rendering our stimuli. We created each shape using a triangular mesh model of a sphere containing over 300,000 triangles. To create variations in the surface, we randomly selected vertices from the model and increased or decreased their distances from the sphere’s origin. We modified the locations of the neighboring vertices in a similar manner to keep the surface smooth. We shaded each object using a local Lambertian shading model without specularities or cast shadows. See Figure 1 for an example. See the Lighting section below for more information on the lighting conditions.




Figure 4: Individual measurements of the angular difference (deg) between the actual and perceived surface normal as a function of light angle (deg). Each line represents data from one subject. The error bars represent the standard error of the mean. Although there are slight variations between subjects, each individual performed optimally when the light was slanted between 11 and 44 degrees above the viewing direction.

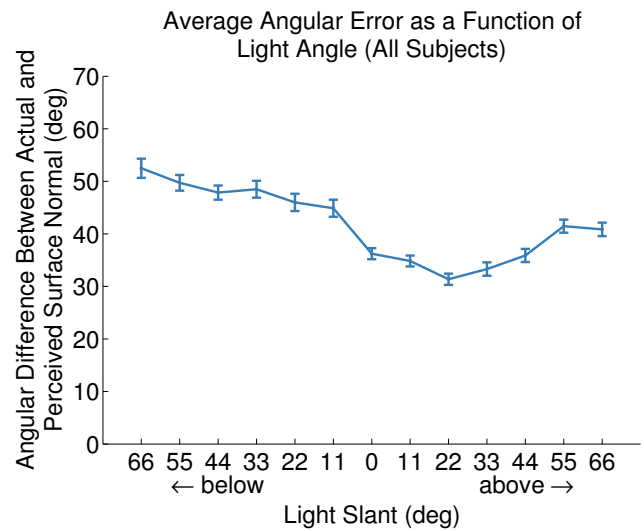


Figure 5: Average angular difference (deg) between the indicated and actual surface normal as a function of light angle (deg). The data depicted in this chart are the combined average for all three subjects after normalizing the data. The error bars represent the standard error of the mean. A light angle of 0 degrees corresponds to an illumination direction from the viewpoint. Errors are highest when the light is positioned below the object. Subjects performed best when the light was slanted 22 degrees above the viewing direction.

3.4 Gauge figure

We instructed subjects to indicate the surface normal at specified points on the surface using a gauge figure [Koenderink et al. 1992]. The gauge figure consisted of a single line with a circle at its base. The base appeared fixed to the surface (Figure 3). We presented the figure using perspective projection to provide additional information about its orientation. The circle was included as an additional cue to orientation. The line and circle were rendered without shading to avoid interfering with the shading of the surface. The orientation of the gauge figure was set randomly at the beginning of each trial. We rendered the gauge figure at a constant size so changes in its size would not reveal information about variations in depth along the surface. The task was conducted with monocular viewing to avoid a binocular cue to flatness that would have otherwise occurred.

3.5 Lighting

For each trial, we illuminated the object using a single directional light source. Because our goal was to study how the light slant (angle between light direction and the viewing direction) affects shape perception, we varied the direction of light such that its slant ranged between 0 and 66 deg in 11-deg increments over the course of the experiment. There is ample research suggesting a light-from-above prior, and a review of relevant literature revealed that the average assumed light source is above and to the left 12 deg from vertical [Sun and Perona 1998; Mamassian and Goutcher 2001; Adams et al. 2004; Champion and Adams 2007]. We therefore constrained our light source to always be tilted in this direction, rotating it to be slanted either above or below the viewing direction (Figure 2). In our coordinate system, this is equivalent to varying the light slant and keeping the light tilt fixed at 102 deg.

3.6 Experimental Design

We varied three parameters throughout the experiment: lighting direction, surface slant, and surface tilt. We predetermined the parameters of each trial to guarantee that a uniform sampling of surface slants and tilts would be tested under each lighting direction. We presented these trials to the subject in a randomized order. The angle between the light direction and the view direction varied from 0 to 66 deg in 11 deg increments, both above and below the viewpoint. We tested surface slants from 0 to 80 deg in 20 deg increments, and surface tilts from -120 deg to 180 deg in 60 deg increments. These parameters yield 390 unique combinations, each of which was tested three times for a total of 1170 trials.

3.7 Procedure

For each trial, we presented the image of a surface to the subject with the gauge figure superimposed at the location to be tested. Using the mouse, subjects oriented the gauge figure until it appeared normal to the surface. The setting was made by clicking the right mouse button. Between trials, a gray screen was displayed for 500ms to minimize the after-image effect on the perceived shading of the next trial. The entire experiment was self-paced, and the data was collected over several one-hour sessions which were broken up into 20 minute blocks. See Figure 3 for a screenshot from the experiment.

4 Results

We recorded the surface location, the setting made by the subject, and the actual surface normal for each trial of the experiment. From these data, we calculated slant and tilt components for both the surface and the gauge figure. If the visual system assumes the direction of illumination is slanted away from the viewing direction at




Figure 6: Average slant error (deg) between the estimated and actual surface normal as a function of light angle (deg). Errors were minimized when the light was slanted approximately 0-22 degrees above the viewpoint.




Figure 7: Average tilt error (deg) between the estimated and actual surface normal as a function of light angle (deg). Similar to other measurements of error, subjects performed optimally when the illumination source was slanted 22 degrees above the viewpoint.

a particular angle, we expect the experimental settings to be most accurate when the lighting condition matches this assumption.

We first assessed subject performance by calculating the angular difference between the subject setting and the actual surface normal. Individual results for the three subjects are shown in Figure 4. The results averaged across the subjects after normalizing the data are shown in Figure 5. These data confirm the light-from-above prior. As the light direction moves from below the viewpoint to above it, average errors decrease. Errors are smallest when the light is slanted 20-30 deg above the viewing direction.

We also analyzed errors in surface slant and tilt estimation as a function of light direction. (Figures 6 and 7). Slant error was calculated as the angular difference (in degrees) between the actual surface slant and the slant indicated by the subject's setting. Tilt error was similarly calculated. These results are consistent with the previous findings, confirming a preference for light from above as well as showing that errors were smallest when the light was slanted 20-30 deg above the viewing direction.

Although these three measures of error are consistent with each other, we consider the angular difference between the actual surface normal and the subject's setting to be a better estimate of performance. Slant and tilt errors can sometimes be misleading. When the surface slant is near zero, small errors in setting the gauge figure can disproportionately contribute to large errors in tilt. Likewise, large errors in setting the gauge figure may not be reflected by slant alone if the error is primarily in the tilt direction.

These findings provide further evidence for a light-from-above prior. Regardless of how we measured the error, subjects performed best when the light direction was above the viewpoint. Furthermore, these results show that the most accurate settings were made when the angle between the light direction and the view direction was 20-30 deg. The results of our study suggest that the visual system may assume light is slanted at this angle when presented with Lambertian shading without cast shadows.

5 Error Simulation

In most perception studies, it is impossible to verify that the results accurately reflect observer percepts. The experimental technique may introduce errors that affect the subject responses. In our case, it is possible that subjects misperceived the slant of the gauge figure. Our experiment was not designed to measure this type of error directly. In order to assess whether such errors might have affected our results, we simulated the error by applying a simple error model to our data. Our model of the error is designed to overestimate the magnitude of these errors, and is therefore a conservative model.

Consider the following example: we present a surface slanted at 55 deg but the subject perceives it to be 45 deg. If the experiment does not introduce any additional error, then we would expect the subject's setting to be 45 deg. This setting would be ideal because it would accurately reflect perception. Alternatively, if the subject systematically underestimates the gauge figure's slant, he or she would likely set it to be less than 45 deg. In this case, the experiment would be introducing an error, and the subject's response would no longer be an accurate measure of perception.

We modeled this error as a simple quadratic function (Figure 8). We do not expect subjects to misperceive the gauge figure's slant near 0 deg or near 90 deg, and our model of the error reflects this assumption. We simulated both a 10 deg and 20 deg maximum underestimation of the gauge figure slant. In both cases, the simulated data produce slightly different results (Figure 9), but the general findings remain the same. Errors in surface perception are minimized when the light is slanted 20-30 deg above the viewpoint. This simulation demonstrates that a systematic underestimation of the gauge figure slant would most likely not affect our findings.

6 Discussion

We demonstrated experimentally that viewers make more accurate estimations of surface orientation for shaded images when the light direction is slanted approximately 20-30 degrees above the viewing direction. These results provide evidence that the visual system assumes this lighting direction in order to disambiguate the shape

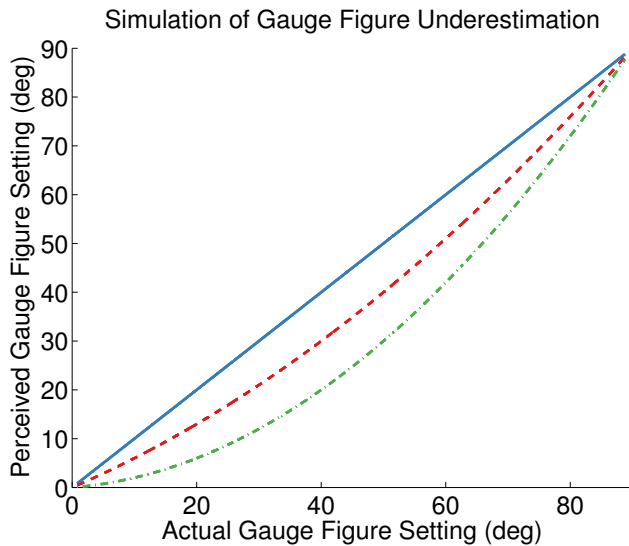


Figure 8: Proposed model of gauge figure misperception. The solid blue line represents a veridical gauge figure model in which the setting perfectly represents the subject's perception. It is possible that subjects underestimated the slant of the gauge figure when it was greater than slant 0 and less than slant 90. We modeled both a 10-deg maximum underestimation (red line) and a 20-deg maximum underestimation (green line). We applied these functions to the data in order to assess how this misperception might have changed our results.

depicted in images rendered using local Lambertian shading without cast shadows.

One of the goals in computer graphics is to convey 3D information on a flat display screen. For a 2D image plane, there are non-pictorial cues to depth that specify a different shape than the effect simulated by shading. Previous research has demonstrated that the brain combines different cues to depth in an optimal fashion in order to achieve an overall depth estimate [Hillis et al. 2004]. In addition to shading, the visual system utilizes cues such as disparity, defocus, and motion parallax to help discern the 3D shape of an object [Watt et al. 2005; Hoffman et al. 2008]. When viewing objects in the real world, all of these cues provide a consistent 3D shape, yet this is not the case with simulated 3D scenes on a 2D display. In fact, the non-pictorial cues in these situations specify that the observer is really viewing a flat surface. Thus, shape cues such as shading need to be effective enough to overcome the flatness cues and thereby yield the desired percept for the simulated object.

Regardless of how we measured subject error, performance was always found to be best when the light was approximately 20-30 deg above the viewpoint. At higher angles, errors increased even though these conditions were still consistent with a light-from-above prior. One explanation is that these extreme light directions illuminated the object with a more raking angle, and a greater portion of the surface was oriented more than 90 deg away from the light source. These surface areas correspond to attached shadows, and thus are completely dark. This shading is due to the local diffuse shading model (Eq. 1). Without adequate variation in the shading of the surface around the gauge figure, it may be impossible to make any estimation of the surface orientation.

To assess how this situation may have affected the results, we ran a second analysis of the data excluding trials in which the gauge figure was positioned in an area of attached shadow (Figure 10). We

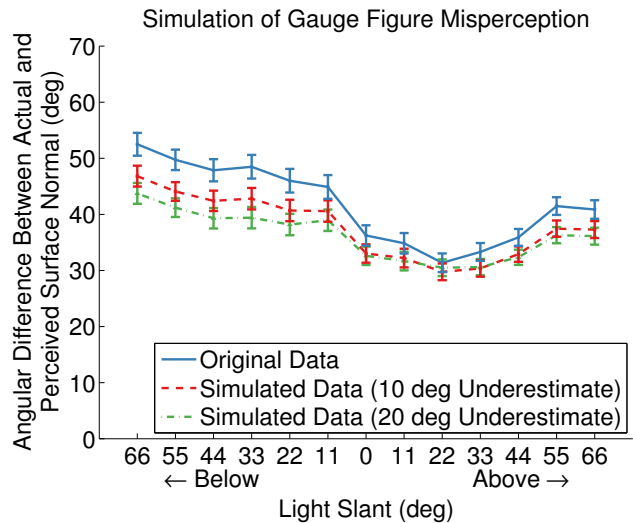


Figure 9: Simulated effect that a systematic underestimation of the gauge figure slant would have on shape perception. These data represent the difference between the actual and perceived surface normal. Although the error introduced by the misperception of the gauge figure changes the data slightly, the overall results remain the same. We conclude that an underestimation of the gauge figure slant would not affect our findings.

excluded 164 of the 1170 trials from each subject (14%). In general, we found a similar pattern of errors in this analysis. The main difference is that errors were slightly smaller when the light was directed from the highest angles above the viewpoint. Based on this analysis, it is unlikely that the overall findings were significantly affected by trials located in regions of attached shadows.

There is also the possibility that cast shadows would affect our results. We used a local Lambertian shading model which does not include computations for cast shadows. The visual system may make incorrect surface estimations when cast shadows would otherwise exist. As such, our results primarily apply to local diffuse shading models that do not compute cast shadows. Despite the simplicity of this model, it is commonly used in computer graphics applications.

As previously mentioned, there is evidence that the visual system has a prior for global convexity [Langer and Bulthoff 2001], which could affect shape perception. Although we do not control for this prior in our experimental design, it is unlikely to affect our findings. Because our stimuli are globally convex, any effect of the convexity prior would most likely improve shape perception and weaken our findings. The fact that we still find an effect of light direction despite this prior only strengthens our results.

It is also possible that our class of surfaces influenced the overall findings. Depending on the geometry of the surface, certain lighting directions may provide more useful shading information compared to other lighting directions. These differences may not be the same for all of the surfaces used in our study, which may affect subject performance. Figure 11 provides example images of different surface types after applying the results from our study. Previous research has attempted to quantify the amount of shading information in an image [Shackled and Lischinski 2001; Gumhold 2002], but it is still not known how the visual system ultimately makes use of this information. Further work is needed to distinguish the affect of surface geometry from the visual system's assumptions about lighting.




Figure 10: Average errors after removing trials in which the gauge figure was in attached shadow. This graph shows the angular difference (deg) between the indicated and actual surface normal as a function of light angle (deg). We excluded 14% of the trials for each subject. We remove these trials from the analysis because they often occurred in areas of the surface that lacked local shading variation.

We acknowledge that there are many sources of error in this type of shape perception study. The flat display screen provides conflicting depth cues, the gauge figure may be misperceived, and the shading may produce errors. We do not know how these sources of error combine to affect the recorded observer percepts. In addition, we cannot infer the strength or shape of the error distribution attributed to changes in the light direction alone. This type of characterization of the assumed light angle would require further study.

7 Conclusion

In this study, we have confirmed the visual system assumes light is above when viewing the image of a shaded 3D surface. Additionally, we have demonstrated that the viewer’s perception of shape is more accurate when the angle between the light direction and viewing direction is 20-30 deg above the viewpoint. This experiment provides evidence that the visual system assumes this angle of lighting when presented with local Lambertian shading information without cast shadows.

7.1 Future Work

There are two remaining questions which deserve further study. First, we would like to investigate the extent to which the gauge figure task accurately reflects observer percepts. Some of the error in the experiment may be attributed to the task itself, so it would be useful to measure this effect. One approach would be to conduct a similar experiment using a known, cue-rich stimulus such as a real object. We can assume the subject would accurately perceive the object’s shape. Any errors in the shape perception task could therefore be attributed to the task itself.

A second question to explore is whether more realistic shading, such as Lambertian shading with cast shadows, would affect the findings. As previously noted, subjects may incorrectly estimate surface orientation for regions where cast shadows would normally exist. An investigation of lighting priors may be more accurate if




Figure 11: Example surfaces displayed using two different light directions. The images on the left were shaded with the light at the viewpoint. The images on the right were shaded with the light angled 26 deg above the view vector. We found shape perception was best for the lighting direction used in the images on the right.

the stimuli provide more realistic shading cues.

Acknowledgments

The authors gratefully acknowledge support from NIH, NSF, Microsoft Research, The Okawa Foundation, and the Alfred P. Sloan Foundation.

References

- ADAMS, W., GRAF, E., AND ERNST, M. 2004. Experience can change the light-from-above prior. *Nature Neuroscience* 7, 10, 1057–1058.
- BELHUMEUR, P., KRIEGMAN, D., AND YUILLE, A. 1997. The bas-relief ambiguity. *IEEE Computer Science Society Conference on Computer Vision and Pattern Recognition*, 1060–1066.
- CANIARD, F., AND FLEMING, R. 2007. Distortion in 3d shape estimation with changes in illumination. *Proceedings of the 4th symposium on Applied perception in graphics and visualization*, 99–105.
- CHAMPION, R. A., AND ADAMS, W. J. 2007. Modification of the convexity prior but not the light-from-above prior in visual search with shaded objects. *Journal of Vision* 7, 13 (10), 1–10.

- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics* 1, 1, 7–24.
- DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive contours for conveying shape. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 848–855.
- GUMHOLD, S. 2002. Maximum entropy light source placement. In *VIS '02: Proceedings of the conference on Visualization '02*, IEEE Computer Society, Washington, DC, USA.
- HILLIS, J., WATT, S., LANDY, M., AND BANKS, M. 2004. Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision* 4, 12 (Dec.), 967–992.
- HOFFMAN, D. M., GIRSHICK, A. R., AKELEY, K., AND BANKS, M. S. 2008. Vergenceaccommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision* 8, 3 (3), 1–30.
- KOENDERINK, J., VAN DOORN, A., AND KAPPERS, A. 1992. Surface perception in pictures. *Perception and Psychophysics* 5, 52, 487–496.
- KOENDERINK, J., VAN DOORN, A., AND CHRISTOU, C. 1996. Perturbation study of shading in pictures. *Perception* 25, 1009–1026.
- LANGER, M., AND BULTHOFF, H. 2001. A prior for global convexity in local shape-from-shading. *Perception* 30, 403–410.
- LEE, C. H., AND HAO, X. 2006. Geometry-dependent lighting. *IEEE Transactions on Visualization and Computer Graphics* 12, 2, 197–207. Member-Amitabh Varshney.
- MAMASSIAN, P., AND GOUTCHER, R. 2001. Prior knowledge on the illumination position. *Cognition* 81, B1–B9.
- MAMASSIAN, P., AND KERSTEN, D. 1995. Illumination, shading and the perception of local orientation. *Vision Research* 36, 15, 2351–2367.
- OLIENSIS, J. 1991. Shape from shading as a partially well-constrained problem. *Computer Vision, Graphics, and Image Processing* 54, 2, 163–183.
- OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert’s reflectance model. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 239–246.
- PHONG, B. T. 1975. Illumination for computer generated pictures. *Communications of the ACM* 18, 6, 311–317.
- PRADOS, E., AND FAUGERAS, O. 2005. Shape From Shading: a well-posed problem? In *Proceedings of CVPR'05*, vol. 2, 870–877.
- RUSINKIEWICZ, S., BURNS, M., AND DECARLO, D. 2006. Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July).
- SHACKED, R., AND LISCHINSKI, D. 2001. Automatic lighting design using a perceptual quality metric. In *EG 2001 Proceedings*, vol. 20(3), 215–226.
- STEVENS, K. A. 1983. Slant-tilt: The visual encoding of surface orientation. *Biological Cybernetics* 46, 183–195.
- SUN, J., AND PERONA, P. 1998. Where is the sun? *Nature Neuroscience* 1, 3 (July), 183–184.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for the off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9, 1105–1144.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2, 265–272.
- WATT, S. J., AKELEY, K., ERNST, M. O., AND BANKS, M. S. 2005. Focus cues affect perceived depth. *Journal of Vision* 5, 10 (Dec.), 834–862.