

Three-dimensional object shape from shading and contour disparities

The School of Psychology, University of St Andrews,
St Andrews, United Kingdom, &
Department of Physics of Man,
Universiteit Utrecht/Helmholtz Instituut,
Utrecht, The Netherlands



Harold T. Nefs

Both non-Lambertian shading, specularities in particular, and occluding contours have ill-matched binocular disparities. For example, the disparities of specularities depend not only on a surface's position but also on its curvature. Shading and contour disparities do in general not specify a point on the surface. I investigated how shading and contours contribute to perceived shape in stereoscopic viewing. Observers adjusted surface attitude probes on a globular object. In [Experiment 1](#), the object was either Lambertian or Lambertian with added specularities. In the next experiment, I removed the Lambertian part of the shading. In [Experiment 3](#), I reduced the disparity of the contour to zero, and in [Experiment 4](#), I removed both cues. There was little effect of shading condition in [Experiment 1](#). Removing the Lambertian shading in [Experiment 2](#) rendered the sign of the surface ambiguous (convex/concave) although all surfaces were perceived as curved. Results in [Experiment 3](#) were similar to those in [Experiment 1](#). Removing both cues in [Experiment 4](#) made all surfaces appear flat for three observers and convex for one observer. I conclude that in the absence of Lambertian shading, observers have categorically different perceptions of the surface depending on whether disparate specular highlights and disparate contours are present or not.

Keywords: stereovision, object shape, 3D shape, specular highlights, contours

Citation: Nefs, H. T. (2008). Three-dimensional object shape from shading and contour disparities. *Journal of Vision*, 8(11):11, 1–16, <http://journalofvision.org/8/11/11/>, doi:10.1167/8.11.11.

Introduction

I asked the question how smoothly shaded and glossy objects are perceived in stereoscopic vision. Because the two eyes look at objects from slightly different vantage points, the images on the two retinas are slightly different. This provides the visual system with cues that are not available in monocular vision alone. It is tempting to regard the disparities of the shading pattern,¹ of specular highlights in particular, and the disparities of the occluding contour of an object as potentially useful cues for three-dimensional shape recovery. Occluding contours, for example, are usually clearly identifiable with sharp boundaries in the two half-images. However, both disparate shading and disparate contours have their limitations and are in general ambiguous cues to the three-dimensional shape of the object. That is, these disparities do not correspond to the position the object occupies in space but depend also on other factors such as the surface's curvature. The conventional shape-from-disparity algorithms do therefore not apply and may in fact give erroneous results. Despite the vast amount of knowledge on stereoscopic shape perception of objects defined by texture, the way the visual system extracts shape information from smoothly shaded and glossy objects in stereoscopic viewing is less advanced. In this

paper, I investigate the separate and combined effects of Lambertian shading, specular highlights, and the occluding contour in stereoscopic shape perception.

Points with the same luminance in the two half-images of a stereoscopically viewed object are only derived from the same point on an object if the Bidirectional Reflection Distribution Function, the BRDF, is Lambertian. In Lambertian reflectance, incident light is scattered in all directions to the same degree (see also Nefs, Koenderink, & Kappers, 2006). It therefore follows that no matter from which angle a Lambertian point is looked at, it will have the same luminance. This is not so for non-Lambertian shading: In general, the luminance of a point is also dependent on the viewing direction relative to the surface. That is, points with the same luminance that seem to match in the two half-images may in reality be projected from different points on the object. Therefore, I say that the disparity of non-Lambertian shading is ill-matched. An example of a BRDF that clearly results in ill-matched disparities is specular reflection. As illustrated in [Figure 1](#), the positions of the specular highlights on the object are different for each eye. According to the disparity, the specular highlights are positioned in front of the object's surface for concave areas and behind the surface for convex areas. In the case of a hyperbolic surface, that is saddle-shaped, the disparities signal a position in front or behind the surface depending on the orientation of the




Figure 1. The positions of specular highlights on the object are different for the two eyes. Specular highlights (Sp) can be regarded as if they were created in front of or behind the surface, depending on the curvature of the surface. The curved line is a cross-section of specular surface. Incident light is coming from the bottom-right of the page. Green lines are the reflections to the left eye and the red lines are the reflections to the right eye. The horizontal direction with respect to the observer is indicated by x and the depth dimension is in the direction of z . A specular highlight is seen where the green and the red lines intersect. The red and green lines intersect in front of the surface when it is concave and behind the surface where it is convex.

principal curvatures of the surface relative to the interocular axis. Because highlights are not infinitely small but spread to some extent over the object's surface, there are usually not only disparities in position of the highlights between the two eyes but also disparities in shape of the highlights. In order not to confound these two disparities, I will only consider the effects of positional disparity of specular highlights between the two eyes by using small specular highlights with sharp boundaries that only have minimal shape disparities (that is, in technical terms, the specular highlights have a high specular exponent). Since the reflection properties and illumination properties are a priori unknown to the observer, it is not clear how the observer should interpret any shading disparity.

There are indications that the disparity of shading of low spatial frequency does not influence shape perception much in stereoscopic viewing; that is, it is a weak cue (e.g., Vuong, Domini, & Caudek, 2006). For example, Arndt, Mallot, and Bülthoff (1995) reported that pseudoscopic viewing of shaded spheres alters the perceived distance to the sphere. In pseudoscopic viewing, the two left and right half images are swapped. Pseudoscopic viewing also makes the sphere appear a bit flatter, but it does not change spheres into concave bowls. Apparently, monocular cues together with a prior for convexity (Hill & Bruce, 1993, 1994) are the dominant factors in these cases and not binocular disparity. As a historical side note,

Wheatstone (1852) already described the failure of turning convex objects into concave ones during pseudoscopic viewing. The specular component of the shading pattern on the other hand is potentially informative in stereoscopic viewing. Blake and Bülthoff (1990, 1991) and Zisserman, Giblin, and Blake (1989) argued that in stereovision convex/concave ambiguities can be resolved by evaluating the 3D position of specular highlights. Blake and Bülthoff (1990, 1991) showed that observers who were looking at convex or concave surfaces defined by disparity, shading, and texture were able to indicate where specularities should occur in 3D space. One of the most important problems with the disparities of specular highlights is that it is relatively easy to understand where they will be located in three-dimensional space given a certain surface shape, but it is far more complicated and probably ambiguous to find the inverse: What is the surface shape given a certain disparity of a specular highlight?

There seems to be consensus in the literature that shaded objects are perceived more veridically when viewed in stereo. Norman, Todd, and Phillips (1995) showed that correlations between “actual” tilt and tilt component of gauge figure settings for globular objects containing just Lambertian shading or specular highlights, or any combination of these cues improved when viewed in stereo. Although R^2 's were low, the correlation of the slant components of the gauge figure settings was also improved when the objects were viewed in stereo. Todd, Norman, Koenderink, and Kappers (1997) also reported better performance on shape tasks for stereoscopic viewing of specular objects as well as a scaling effect on the perceived shape when specular highlights were present. Bülthoff and Mallot (1988) found that more observers perceived more depth when disparate shading was used than when identical shading was used for both eyes.

Obviously, all monocular shading cues that are available in monoscopic viewing are available as well in stereoscopic viewing. It is well known that Lambertian shading is an important cue for shape perception in monoscopic viewing. It has also been investigated how the perceived shape changes with changes in the shading pattern (Nefs, Koenderink, & Kappers, 2005). The effects of specular highlights on perceived shape in monocular viewing conditions are, however, not fully understood yet. Specularities carry a potential amount of information on the surface shape in monoscopic vision (e.g., Flemming, Torralba, & Adelson, 2004; Longuet-Higgins, 1960; Nefs et al., 2006; Oren & Nayar, 1996). However, previous reports on the actual use of specular highlights for shape perception appear contradictory. Todd and Mingolla (1983) reported that glossiness enhances the perception of curvature of cylindrical surfaces. Mingolla and Todd (1986) did not find any significant effects of specular highlights on the perceived shape of ellipsoids. Norman et al. (1995) reported that specular highlights on otherwise black objects give a compelling sense of 3D shape on their own. Nefs et al. (2006) were not able to find significant

scaling or shearing effects on the perceived shape for the presence of specular highlights on otherwise matte globular objects. Norman, Todd, and Orban (2004) showed that specular highlights can provide important cues to distinguish the 3D shape of objects.

Another binocular cue that may be of use for three-dimensional shape perception is the disparity of the occluding contour. However, the disparity of the occluding contour is also ill-matched in stereoscopic viewing (e.g., Anderson & Nakayama, 1994; Norman & Raines, 2002). The contour generator is the curve over an object where the object curves to occlude itself from sight. The contour generator is different for both eyes. Therefore, the horizontal shift of the contour between the two half-images is not only determined by the binocular parallax of the object's surface, but also because the contour generators for the left and right eye are different. Contour disparity is not unique for a specific three-dimensional shape because there are many different pairs of contour generators that give the same disparity. Consider, for example, Figure 2. Any constructed object that touches all

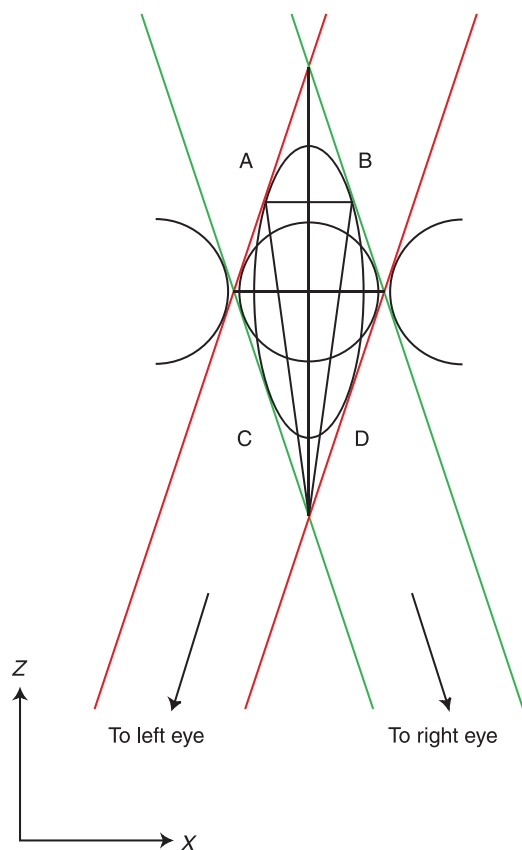


Figure 2. The false target problem for occluding contours. Whatever is tightly enclosed by the lines A, B, C, and D creates the same silhouettes and is thus, as far as the occluding contour is concerned, indistinguishable to the observer. We make the assumption that what we are looking at is an object and not an aperture; the outer two half-circles also give rise to the same disparity between the half-images.

four sides of the bounding volume, whether it is convex, hyperbolic, or concave, voluminous or flat, is a veridical solution for the object's shape.

In binocular viewing, the object must be tightly enclosed within the bounding volume that is constructed with the viewing rays that touch but do not intersect the object. This is a considerable improvement over the case when the object is viewed with one eye only. Information about what is happening inside the bounding volume, as far as the disparity of the contour is concerned, is not available to the observer. The bounding volume is a quite severe restriction and with some further assumptions such as that the object has to have volume and is not a two-dimensional curved sheet the number of possible solutions decreases rapidly. It is, for example, not too complicated to calculate the most-voluminous ellipse that still fits in the bounding volume. At this point one has made several assumptions, e.g., connectivity, has volume, and a 2-symmetry of the shape. Note also that the bounding volume is Euclidean if vergence and version of the eyes are "known" or distance and azimuth to the object are known from other sources, otherwise the bounding volume falls within a projective equivalence class. The situation in Figure 2 is two-dimensional. It describes the occluding points of a flat horizontally oriented shape. In order to generalize to three-dimensional objects, one has to imagine a vertical stack of the 2-dimensional shapes.

In Figure 3, I show the occluding contour of a three-dimensional object, not unlike the one that I used in the Experiments below. Figure 3 is more schematic than the real stimulus to illustrate the ambiguity more clearly. Three specific shapes could be reconstructed based on a disparate occluding contour using different simple filling-in algorithms. The thickness of the contours in Figure 3 indicates the disparity, and hence the depth if interpreted in the conventional way, of the occluding contour. Filling-in the shape with iso-depth lines vertically will result in negatively curved cylindrical shape. Filling-in the occluding contour horizontally with iso-depth lines will lead to a positively curved cylindrical shape. Of course, one can compromise between these previous two solutions and come up with a saddle shape solution. Note that, the first two solutions are limiting cases: Each iso-depth line in Figure 3 would be a straight line between the point where A and C cross and where B and D cross in Figure 2, or a vertical line connecting the points where A and C (or B and D) cross in different horizontal slices of the 3D shape. These two solutions are thus special solutions as the front and back of the reconstructed object touch everywhere. In addition, many other solid objects can be formed. Nevertheless, when measuring perceived shape with stimuli in which the disparity of the occluding contour is the only available cue, one might expect categorically different interpretations of the stimulus.

So far it remains largely to be determined what humans are able to infer from disparate occluding contours. Norman and Raines (2002) showed that ordinal depth




Figure 3. Three possible surface solutions when the contour disparity is given a conventional interpretation. The thickness of the lines indicates the size of disparity, which is interpreted here as distance from the observer in depth. Distance increases with decreasing line thickness. Even though all three figures have the same contour disparity, at least three different interpretations can be given: Negatively curved, positively curved, and saddle shaped.

discriminations improve for objects of which only the silhouette was visible (no shading, no highlights, no texture) when disparate occluding contours were used compared to the same objects viewed monocularly. Todd et al. (1997) found that observers' surface attitude settings on Lambertian surfaces were more reliable for stereoscopic viewing than for monocular viewing of matte, glossy, or textured objects. Norman et al. (2004) using black objects of which only the contour was visible did not find significant differences when single images of the objects were shown to observers or when multiple images (i.e., results were collapsed over stereoscopic presentation conditions and motion conditions) were used.

As for shading, the shape of the occluding contour is also a monocular cue. In monoscopic viewing, the occluding contour is found to be highly important for three-dimensional shape perception. When the occluding contour is made invisible the perceived shape is much affected. For example, Erens, Kappers, and Koenderink (1993) showed that removing the occluding contour rendered the sign of the surface curvature ambiguous for Gaussian shapes. Some authors take the fair view that the occluding contour may provide sufficient information for three-dimensional shape perception if accidental, unlikely, situations are excluded (e.g., Tse, 2002).

The purposes one may use shading and contour disparities for are thus limited because of the ambiguities of these cues for 3D shape but do exist and depend on how restricted the visual system is by the assumptions it makes. In order to find out what the separate and combined effects of Lambertian shading, specular highlights, and the occluding contours are in stereoscopic vision, I conducted four related experiments. In all four experiments, I used stereo images of the same globular object. A substantial part of the surface was saddle-shaped such that the binocular disparities of specular highlights change when the object is rotated around the line of sight. I hypothesized first that if specular highlights were present, the perceived depth in the object would increase. Further, I hypothesized that if the hyperbolic part were orientated such that the surface is convex in horizontal direction and concave in vertical direction, the specular highlight would be in front of the surface; if the surface is rotated 90 degrees around the line of sight, the specular highlight is behind the surface. If the observer does not correctly interpret the binocular disparities of specular highlights but interprets them as part of the surface instead, the perceived shape of the object would change when the object is rotated around the line of sight. I presented the object in three different rotation angles around the line of sight. In [Experiment 1](#), the object could either be reflecting in a Lambertian fashion or be Lambertian with added specular highlights. Highlights were clearly identifiable by color and luminance contrast and had reasonably sharp boundaries. The Lambertian shading component changes in both reflection conditions in the same way. In the specular condition, the specular highlights have different shifts between the two eyes for each rotation angle. Because of the asymmetry of the object the disparity of the occluding contour also varies with the rotation around the line of sight. In the next three experiments, I used the same experimental design as in [Experiment 1](#) but I systematically eliminated two cues, namely, Lambertian shading and contour disparity. In [Experiment 2](#), the objects were completely black or black with added specular highlights. In [Experiment 3](#), the contour disparity was set to zero everywhere along the contour. I did this by cutting about 5 percent from the shaded area that was nearest to the occluding contour. I used the contour of the shape as seen from the cyclopean

viewpoint as a mask. The resulting shape of the contour was thus the same for the two eyes. In the fourth experiment, the objects were either black or black with added specular highlights and their contour disparity was set to zero everywhere along the contour.

Experiment 1: All cues

Method

Observers

Four observers, three male and one female, participated in this experiment. All participants in this study were students at Utrecht University between 18 and 30 years of age. Observers participated in only one experiment in this project. They were naive to the purposes of the experiment. None of them had seen any of the stimuli prior to the experiment. I pre-screened the observers for their visual abilities. All observers had normal or corrected-to-normal far and near acuity, stereovision as determined by the TNO stereo-test (1972), and accurate color vision, as determined with the Ishihara plates (1989).

Procedure

Observers were seated in front of a computer screen in a dimly lit room. The only direct light source in the room was the computer's monitor. The distance from the screen to the observer was 75 cm. The stereo images were presented next to each other on the screen. The observers viewed the stereo images through a standard dual-mirror stereoscope.

A small red probe, which consisted of a circle and a small line, was superimposed on the left half-image. The probe resembled the shape of a pushpin. The observers could alter the aspect ratio and orientation of the probe with simple movements of the computer mouse. The observers were instructed to manipulate the probe such that it seemed to be a circle painted on the pictorial surface of the displayed object with an outward surface normal. Observers find this a quite natural task and experience no difficulty in performing it. This method has been used successfully many times before for monocular viewing (e.g., Koenderink, van Doorn, & Kappers, 1992; Nefs et al., 2005, 2006), as well as for stereo viewing (e.g., Doorschot, Kappers, & Koenderink, 2001; Todd et al., 1997). I identified 149 different probe positions within the occluding contour at regularly spaced intervals. The probe was placed four times at each of these positions.

Each experiment was divided over four sessions that were separated by at least a day. In each session every image pair was sampled once completely; a new pair of stereo pictures was presented after all probe positions had been set. Typically, only a few seconds were needed for a single setting of the probe. A session lasted about 45–60 minutes each (6 images \times 149 probe positions =

894 settings). The order of the images and the order of the probe positions were randomized anew for each session and each observer.

Stimuli

I made six stereo images of a globular convex object. I created the object by modulating the radius of a unit sphere as a function of azimuth and elevation. The object was oriented such that hyperbolic, that is saddle-shaped, and convex areas can both be seen in the picture.

The displayed objects were built from 20,480 triangular faces with interpolated RGB-values; as a result, the displayed objects looked very smooth. The six stereo pairs of the object that are used in this study are shown in Figure 4. The right half image is taken from the object as seen from 2.27 degrees to the right of the cyclopean eye, and the left image is taken from the object as seen from 2.27 degrees to the left of the cyclopean eye with the center of the object as the axis of rotation. The projection was orthogonal. All observers agreed that all image pairs yielded strong impressions of three-dimensional objects that were floating in front of them.

The BRDF of the object was either Lambertian or Lambertian with added specular highlights. I refer to these BRDF's as "matte" and "glossy," respectively. The matte component of the objects' color was a saturated green and the specular highlight was white. I used the same model as Nefs et al. (2006) for computing the object's luminance as described in Equation 1. This model is a standard OpenGL method for calculating luminance:

$$\begin{aligned}
 L_{\text{Total}} &= L_{\text{Ambient}} + L_{\text{Lambertian}} + L_{\text{Specular}}, \\
 &\text{where} \\
 L_{\text{Ambient}} &= 0.15 \\
 L_{\text{Lambertian}} &= g(\mathbf{n} \cdot \mathbf{s}) \\
 L_{\text{Specular}} &= h \left(\frac{\mathbf{v} + \mathbf{s}}{|\mathbf{v} + \mathbf{s}|} \cdot \mathbf{n} \right)^e.
 \end{aligned} \tag{1}$$

In this equation, the total luminance (L_{Total}) is the linear summation of an ambient (L_{Ambient}), Lambertian ($L_{\text{Lambertian}}$), and specular component (L_{Specular}). In these experiments, the ambient component is set to 15% of the monitor's maximum radiance. The Lambertian component is the vector product of the surface normal \mathbf{n} and the illumination direction \mathbf{s} , multiplied by an attenuation factor g . Here, g is set to 0.8; that is, 80% of the monitor's maximum radiance. The specular component is the dot product of the normalized sum of the viewing direction \mathbf{v} and the illumination direction \mathbf{s} , and the surface normal \mathbf{n} . The specular exponent e is set to 125. For the matte objects, the specular attenuation factor h is set to zero, and for the glossy objects h is set to 1, that is 100% of the monitor's maximum radiance. I simulated a collimated parallel light source. The light was coming from the




Figure 4. The stimuli of [Experiment 1](#). The matte conditions are in the left column and the glossy conditions are in the right column. The images for the rotation angles of 0, 35, and 70 degrees are at the top, middle, and bottom, respectively. The images are arranged in such a way that they allow for both crossed and uncrossed free fusion. The images marked “L” are for the left eye, those marked “R” are for the right eye. The left images contain the attitude probe; its size is proportional to the stimuli that are displayed here.

upper-right region behind the observer and its color was white. The illumination direction was the same relative to the observer for all conditions. I used three different orientations of the object: the object was rotated 0, 35 or 70 degrees anticlockwise around the cyclopean line of sight. The object measured about six degrees of visual angle in each half-image on the computer screen. The background was gray.

Apparatus

I used a MacIntosh PowerPC with dual 1.0-GHz G4 processors and a dual NVIDIA-3D graphics card to display the stereo images on a 22-in. Radius monitor. The vertical refresh rate was set at 85 Hz with a screen resolution of 1280*1024 pixels. Images were rendered by custom-made software that was written with MacIntosh’ implementation of OpenGL 1.2.1 in OS-X.

Analysis

I interpreted the probe settings as depth gradients. With depth gradients, I mean the change in depth with a change in horizontal or vertical direction, $\{\delta z/\delta x, \delta z/\delta y\}$. From a set of depth gradients, it is possible to calculate a best-fitting surface. I describe the algorithm used, along with a numerical example, in the [Appendix A](#).

Depth values are expressed in the same units as the horizontal and vertical screen coordinates. The stereoscopic half-images are spanned each over half of the computer’s screen (640 pixels wide). The left edge of the stereoscopic half-image is assigned the value -1 *screen unit* and the right edge $+1$ *screen unit*. Depending on the experimental condition, the maximum horizontal distance between gauge figure locations was 1.12, 1.10, and 1.0 screen units for the 0, 35, and 70 degree rotation conditions, respectively. The maximum vertical distance between gauge figure locations was 0.96, 0.98, and 1.10 screen units. These horizontal and vertical measures of the depicted objects are needed to form an idea of the width (or height) to depth ratio of the reconstructed surface.

In order to evaluate the similarity between surface reconstructions, I performed several regression analyses. I evaluated the correspondence between two surface solutions with a straight regression analysis as shown in [Equation 2](#) and for an affine regression model as shown in [Equation 3](#). The affine regression model takes the horizontal and vertical positions (x, y) of the gauge figure in the image plane as additional predictor variables.

$$z' = az + d, \quad (2)$$

$$z' = az + bx + cy + d. \quad (3)$$

In [Equations 2](#) and [3](#), a comparison is made between two surface reconstructions. In the analyses that are reported below, these surface reconstructions may belong to different experimental conditions, sessions, or observers. In [Equations 2](#) and [3](#), z and z' are the depth values of the first and the second surface, respectively. The horizontal and vertical coordinates in the image plane are denoted with x and y . A translation in depth is mediated by the constant d ; b and c are the shearing parameters; a scales the surface in depth. R^2 values are invariant for translation and scaling in depth in the straight regression model; in the affine regression model, R^2 values are in addition to invariance for translation and scaling also invariant for shearing transformations.

Results

[Figure 5A](#) shows the interpreted depth gradients for observer A4, averaged over all four sessions. [Figure 5B](#) shows the surface reconstructions for these gradients.

The surface solutions were quite similar for different observers and different sessions; the data in [Figure 5](#) are representative for all observers and all sessions.

I found that the average R^2 's between sessions, when averaged over all combinations of two sessions and over all six conditions, ranged for straight regression from 0.71 to 0.92 for different observers; the mean R^2 across observers was 0.82. The mean standard deviation of the scaling factors between sessions across observers, when settings were averaged over all conditions, was about 18%. This standard deviation is an indication of the amount of scaling between sessions; the mean scaling factor is always close to 1 because for example if session 1 is scaled 0.9 with respect to session 2, session 2 is scaled $1/0.9$ with respect to session 1. Hence, on average close to 1. With the affine regression model average R^2 values improved to a range between 0.89 and 0.96 with the mean R^2 across observers being 0.92. The standard deviation of the scaling factors between sessions was again about 18% with a standard deviation in the shearing factor of $\{0.05, 0.04\}$. The arctangent of the length of a shearing vector converts the vector to an angle, in this

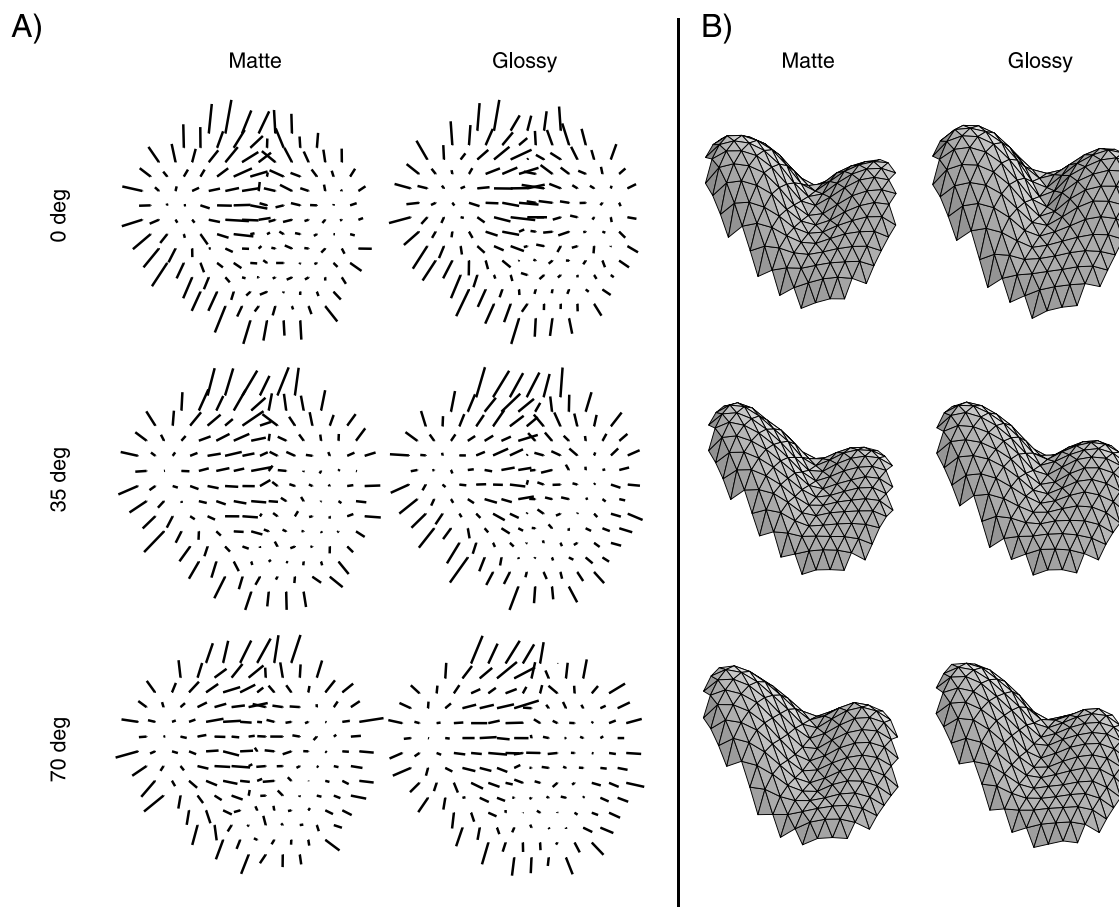


Figure 5. Results for [Experiment 1](#). (A) Gradient settings by observer A4, averaged over all four sessions. The length of each line represents the size of the gradient, that is the steepness of the surface. The direction of each line represents the direction in which the surface changes fastest in depth. The dot at the end of each line represents the position of the surface attitude probe. The contour is taken from the left eye's view. (B) Surfaces constructed from the gradient settings in panel A.

case thus about 3.6 degrees. Since most of the differences seem to be explained by linear transformations, I averaged the depth values over all sessions for each observer in the remainder of these analyses. There was considerable depth in all reconstructed surfaces. The perceived depth, as indicated by the difference between the highest and lowest depth values, ranged from to 0.17 screen units for observer A3 in the “Matte/0 degree” condition to 0.39 screen units for observer A4 in the “Matte/45 degree” condition.

R^2 values for comparisons between observers are lower than between sessions of the same observer. They ranged from 0.45 to 0.85 between different observers for the straight model when averaged over all conditions. The mean R^2 was 0.60. The standard deviation of the scaling factors was 31%. R^2 values ranged between 0.57 and 0.89 with the mean being 0.75 for the affine regression model. The standard deviation of the scaling factors was about 24% and the shearing vector about $\{0.07, 0.02\}$. I did not average probe settings over different observers in subsequent analyses in this section.

Next, I calculated the correlation between the depth values in different BRDF conditions. R^2 's for the correlation between the depth values for different BRDF's ranged from 0.90 to 0.99. With the affine model R^2 values ranged between 0.95 and 0.99. Results, averaged over observers, along with the best estimates of the model parameters are summarized in Table 1. I tested the significance of the best parameter estimates with one-sample T -tests. Scaling parameters that are significantly different from 1, and shearing parameters significantly different from 0 are indicated with asterisks in Table 1. I also calculated the significance of the improvement in R^2 's for the affine regression model over the straight model. Significant improvements in R^2 's are also indicated in Table 1.

R^2 's for the correlations between the depth values for different rotation angles ranged from 0.38 to 0.93. With the affine model R^2 's ranged between 0.58 and 0.98. Results are again summarized in Table 1. In these

comparisons, R^2 's are lower than in the previous analyses. This is related to a shift in the exact position of the hyperbolic area (see Figure 5B).

Experiment 2: No Lambertian shading

Method

Observers

Four new observers, three male one female, participated in this experiment. They did not participate in any of the other experiments in this study. Further particulars are as in Experiment 1.

Procedure and apparatus

As in Experiment 1.

Stimuli

I used the same experimental design as in Experiment 1. However, in the current experiment, the Lambertian component of the shading was zero, that is the object is pitch black, or pitch black with specular highlights. An example of one of the stimuli, namely black with specular highlights at a rotation angle of 35 degrees, is shown in Figure 6A. All stimuli in this experiment are shown in the online materials (Supplement 1).

Results

Surface reconstructions are very different for different observers, and sometimes even for different sessions for the same observer for the black objects without specular

Comparison		Straight model, $z' = az + d$		Affine model, $z' = az + bx + cy + d$			
		R^2	a	R^2	a	b	c
Matte-glossy	0	0.93 (0.01)	0.97 (0.13)	0.97 (0.02)	1.02 (0.08)	0.02 (0.03)	0.02 (0.03)
	35	0.93 (0.02)	0.89 (0.13)	0.97 (0.01)	0.99 (0.13)	0.03 (0.01)	0.01 (0.02)
	70	0.97 (0.02)	1.01 (0.06)	0.98 (0.01)	0.97 (0.09)	-0.01 (0.03)	-0.00 (0.01)
0–35	Matte	0.79 (0.15)	0.93 (0.09)	0.95 (0.04)	0.88** (0.03)	-0.06* (0.03)	-0.05 (0.03)
	Glossy	0.79 (0.06)	0.85 (0.05)	0.94 (0.03)	0.85* (0.06)	-0.04 (0.03)	-0.06* (0.01)
35–70	Matte	0.86 (0.06)	0.97 (0.20)	0.90 (0.07)	0.95 (0.16)	-0.02 (0.02)	0.02 (0.04)
	Glossy	0.83 (0.07)	1.06 (0.20)	0.92 (0.07)	0.95 (0.17)	-0.05 (0.04)	0.01 (0.04)
70–0	Matte	0.67 (0.23)	0.76 (0.19)	0.78 (0.18)	0.94 (0.20)	0.06** (0.01)	0.03 (0.03)
	Glossy	0.57 (0.16)	0.68* (0.13)	0.79 (0.15)	1.01 (0.19)	0.10* (0.06)	0.05 (0.04)

Table 1. R^2 's and best-fit model parameters averaged over all four observers for Experiment 1. Standard deviations are given in brackets. Scaling parameters that are significantly different from 1 and shearing parameters significantly different from 0 are indicated with asterisks. Significant improvements of R^2 's for the affine model over the straight model are also indicated with asterisks. Notes: * $p \leq 0.05$, ** $p \leq 0.01$, two-sided.

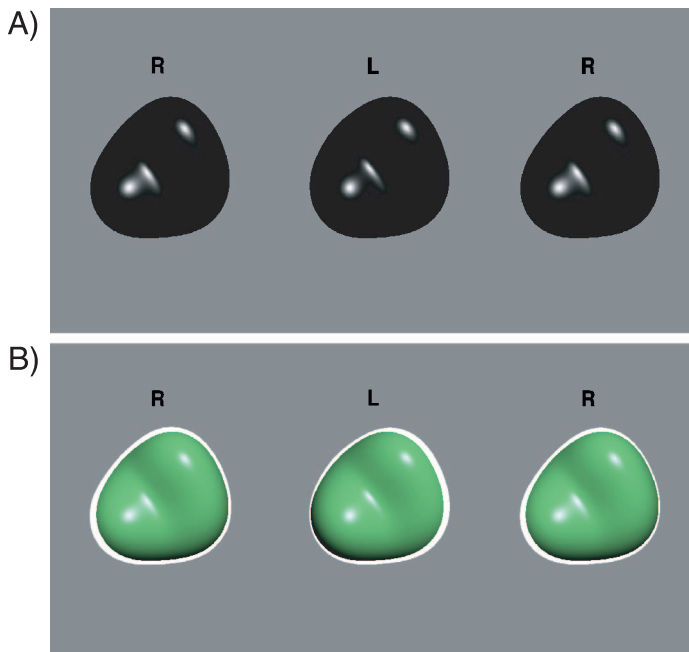


Figure 6. (A) Example of one of the stimuli in [Experiment 2](#): The object has a disparate occluding contour and disparate specular highlights but no Lambertian shading. (B) Example of one of the stimuli of [Experiment 3](#): We clipped a small strip along the contour of about 5% of the area covered by the object in the half-images. We used a 95% scaled instance of the object seen from the cyclopean viewpoint as a mask. Thus, 95% of the Lambertian shading is retained, as well as the specular highlights and the cyclopean shape of the occluding contour, but not contour disparity. The clipped area is shown here in white for illustrative purposes but was shown in the same gray as the background in the actual experiments.

highlights. There was however considerable depth in all reconstructions. The smallest difference between the lowest and highest depth value in the reconstructions was 0.18 screen units for observer B1 in the fourth session for the “Matte/70 degrees” condition. The highest depth range (0.91 screen units) was found for observer B4 in the first session in the “Matte/0 degrees” condition.

The reconstructions for the black objects with specular highlights are more similar between observers and between sessions. However, they do show another peculiar finding. The specular highlights are not perceived as such but rather as bumps on an otherwise uni-dimensionally curved surface. Examples from the reconstructions are presented in [Figure 7](#). The examples in [Figure 7](#) cover all shapes that I have observed in this experiment. The reconstructions seem to be categorically different and not related to any of the experimental factors. Therefore, I do not engage in regression analyses in this section since that is clearly uninformative. In accordance to what I suggested in [Figure 3](#), these results suggest that a disparate contour is not in itself sufficient to lead to an unambiguous percept.

Experiment 3: Zero contour disparity

Method

Observers

Four new observers, three male and one female, participated in this experiment. They did not participate in any of the other experiments in this study. Further particulars are as in [Experiment 1](#).

Procedure and apparatus

As in [Experiment 1](#).

Stimuli

Again I used the same experimental design that was used as in [Experiment 1](#), the three rotation angles (0, 35 and 70 degrees) and two BRDFs (matte and glossy). Illumination conditions were as in [Experiment 1](#). A small strip was cut away near the occluding contour that reduced the disparity to zero everywhere along the contour. I used the silhouette of the same object seen from the cyclopean viewpoint, scaled to 95 percent as a mask. This is illustrated in [Figure 6B](#). In [Figure 6B](#), the strip of shading that was cut from the half-images is shown in white for illustration purposes, but was presented in the average gray of the background in the actual experiment.

Results

The surface solutions are quite similar to each other for different sessions. A representative sample of the data is shown in [Figure 8](#). I proceeded with the same analyses as in [Experiment 1](#). First, I looked at the correspondence between the surface solutions in different sessions for each observer. R^2 's for straight regression between sessions, when averaged over all conditions, ranged between 0.80 and 0.88 with the mean R^2 across observers being 0.81. The standard deviation of the scaling factors, when averaged over all conditions and observers, was about 25%. With the affine regression model, R^2 's improved to a range between 0.86 and 0.93 with the mean R^2 across observers being 0.88. The standard deviation in the scaling factors was about 20% and the standard deviation in the shearing vector was {0.03, 0.03}. Since most of the differences between sessions are explained by linear transformations, I averaged the observer's settings over all sessions for each observer in the remainder of the analyses below. The minimum depth range (0.19 screen units) in the surface reconstructions was for observer C4

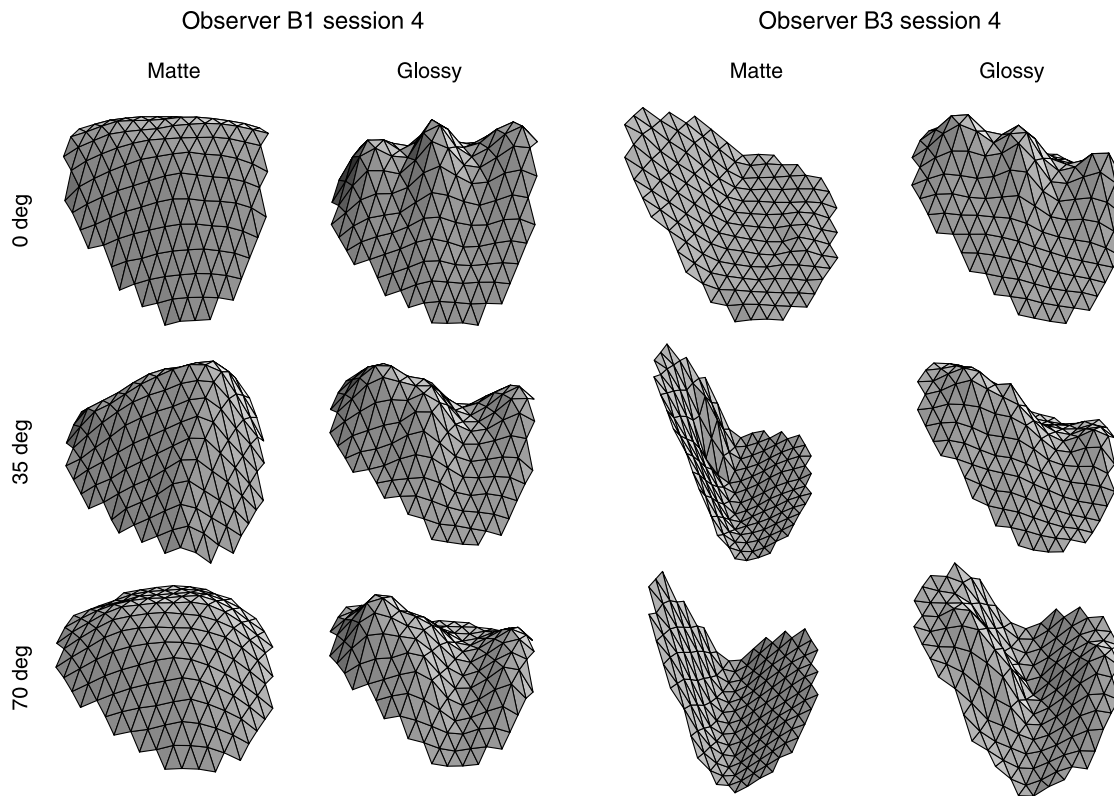


Figure 7. Part of the results for [Experiment 2](#): The surfaces on the left are constructed from the settings made by observer B1 in session 4, and the surfaces on the right are from the settings of observer B3 in session 4. These two observers show the range of different surface reconstructions over the data set.

in the “Matte/0 degrees” condition. Observer C2 had the highest depth range (0.29 screen units) in the “Glossy/0 Degree” condition.

R^2 's between observers, when averaged over all conditions, ranged between 0.45 and 0.85, with the mean being 0.60. The standard deviation of the scaling factor between observers was about 31%. With the affine regression model R^2 's improved to a range between 0.57 and 0.89, with the mean being 0.75. The standard deviation in the scaling factor was about 25% and the standard deviation in the shearing vector was {0.07, 0.02}. I did not average the observers' settings over observers in the remainder of the analyses in this section.

Second, I evaluated the correspondence between the surface solutions for different BRDF conditions. R^2 's for a straight regression model ranged between 0.84 and 0.99 for different observers with the mean being 0.95. With the affine regression model, R^2 's improved to a range between 0.95 and 0.99 with the mean being 0.97. R^2 's averaged over observers as well as the best-estimate model parameters for a straight and an affine regression model are summarized in [Table 2](#). I tested the best-fit model parameters for their significance with one-sample T -tests. Significance levels are indicated in [Table 2](#) with asterisks. I also calculated the significance of the improvement in R^2 for the affine model over the straight

model. Significant improvements of R^2 's are indicated with asterisks in [Table 2](#).

Next, I compared surface solutions for different rotation angles of the displayed object around the line of sight. R^2 's for different comparisons ranged between 0.24 and 0.92 with the mean being 0.75. With the affine regression model, R^2 's improved to range between 0.58 and 0.95 with the mean being 0.85. R^2 's as well as the best-estimate model parameters for a straight and an affine regression model are summarized in [Table 2](#) along with their significance levels.

Experiment 4: No Lambertian shading and zero contour disparity

Method

Observers

Four new observers, one male and three female, participated in this experiment. They did not participate in any of the other experiments in this study. Further particulars were as in [Experiment 1](#).

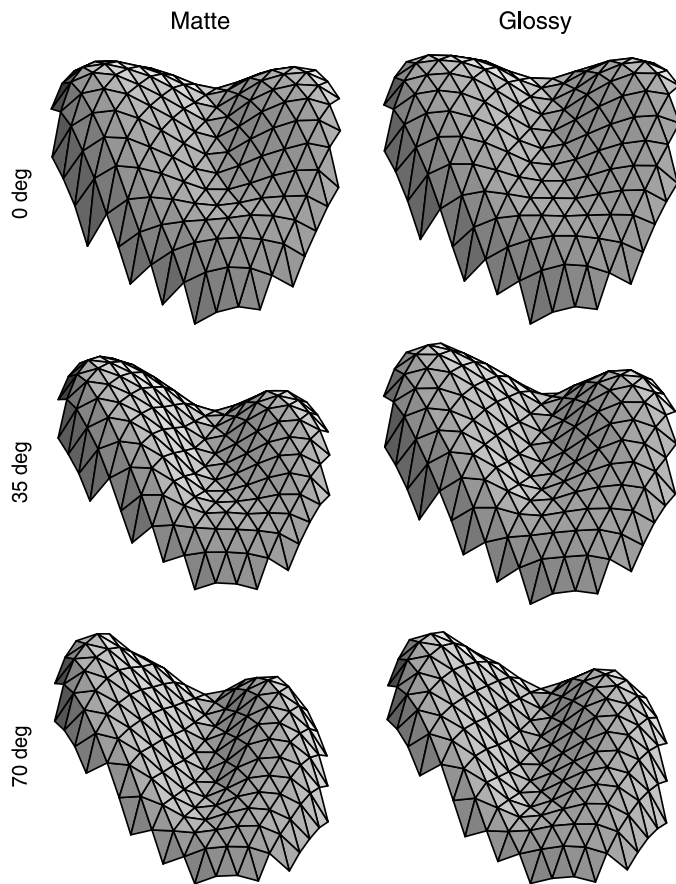


Figure 8. Results for Observer C3 in Experiment 3, averaged over all sessions. These data are representative for all observers.

Procedure and apparatus

As in Experiment 1.

Stimuli

In this Experiment I removed the Lambertian component of the shading, like I did before in Experiment 2. I

also reduced the contour disparity to zero along the contour in the same way as in Experiment 3 as illustrated in Figure 6. All stimuli are shown in the online materials (Supplement 2). The rest is as in Experiment 1.

Results

Results are highly consistent between sessions and between observers D1, D2, and D4 in this experiment. In all instances the surfaces, with or without specular highlights are apparently interpreted as being flat. Observer D4 shows a little bump at the location of the left specular highlight in all three rotation conditions. Observer D2 shows a few extreme gauge figure settings that are not consistent over different sessions and are likely situations where the observer clicked on the computer mouse to soon when adjusting the gauge figure. When the surfaces were averaged over all sessions, of all reconstructions, Observer D4 had the lowest depth range (0.001 screen units) in the “Matte/35 degrees” condition. In spite of the little bump on the left specular highlight, in no condition did the depth range of the surface reconstruction exceed 0.06 screen units for this observer. Observer D2 (due to only a few extreme settings) had a maximum depth range (0.13 screen units) in the “Glossy/35 degrees” condition when averaged over all four sessions. As can be seen in Figure 9, there are only small variations over the surface solutions that are due to imprecision in the settings. Figure 9A is taken from observer D2 but is representative for all sessions and observers D1 and D4. There does not seem to be a need for further regression analyses in this experiment as there does not appear to be any systematic pattern. R^2 's would be very low and non-informative because the surfaces are mostly flat.

Observer D3 showed a different response pattern than the other three observers. She saw all surfaces as

Comparison		Straight model, $z' = az + d$		Affine model, $z' = az + bx + cy + d$			
		R^2	a	R^2	a	b	c
Matte-glossy	0	0.97 (0.01)	1.08 (0.10)	0.97 (0.01)	1.06 (0.08)	0.00 (0.01)	0.01 (0.01)
	35	0.93 (0.06)	0.93 (0.15)	0.97 (0.02)	0.99 (0.07)	0.01 (0.04)	0.00 (0.03)
	70	0.94 (0.04)	0.97 (0.11)	0.97 (0.01)	0.98 (0.07)	-0.01 (0.02)	0.01 (0.01)
0–35	Matte	0.78 (0.20)	0.88 (0.09)	0.94 (0.02)	0.86** (0.04)	-0.03 (0.07)	-0.03 (0.06)
	Glossy	0.85 (0.09)	0.81* (0.08)	0.93 (0.02)	0.81** (0.06)	-0.02 (0.03)	-0.03 (0.03)
35–70	Matte	0.83 (0.03)	0.88 (0.14)	0.87 (0.02)	0.83* (0.11)	-0.01 (0.04)	0.00 (0.02)
	Glossy	0.81 (0.11)	0.90 (0.11)	0.89 (0.01)	0.85* (0.06)	-0.03 (0.05)	0.01 (0.02)
70–0	Matte	0.61 (0.24)	0.84 (0.36)	0.74 (0.10)	1.00 (0.15)	0.02 (0.09)	0.03 (0.05)
	Glossy	0.62 (0.26)	0.93 (0.39)	0.73 (0.11)	1.06 (0.18)	0.04 (0.06)	0.03 (0.04)

Table 2. R^2 's and best-fit model parameters averaged over all four observers for Experiment 3. Standard deviations are given in brackets. Scaling parameters that are significantly different from 1 and shearing parameters significantly different from 0 are indicated with asterisks. Significant improvements of R^2 's for the affine model over the straight model are also indicated with asterisks. Notes: * $p \leq 0.05$, ** $p \leq 0.01$, two-sided.

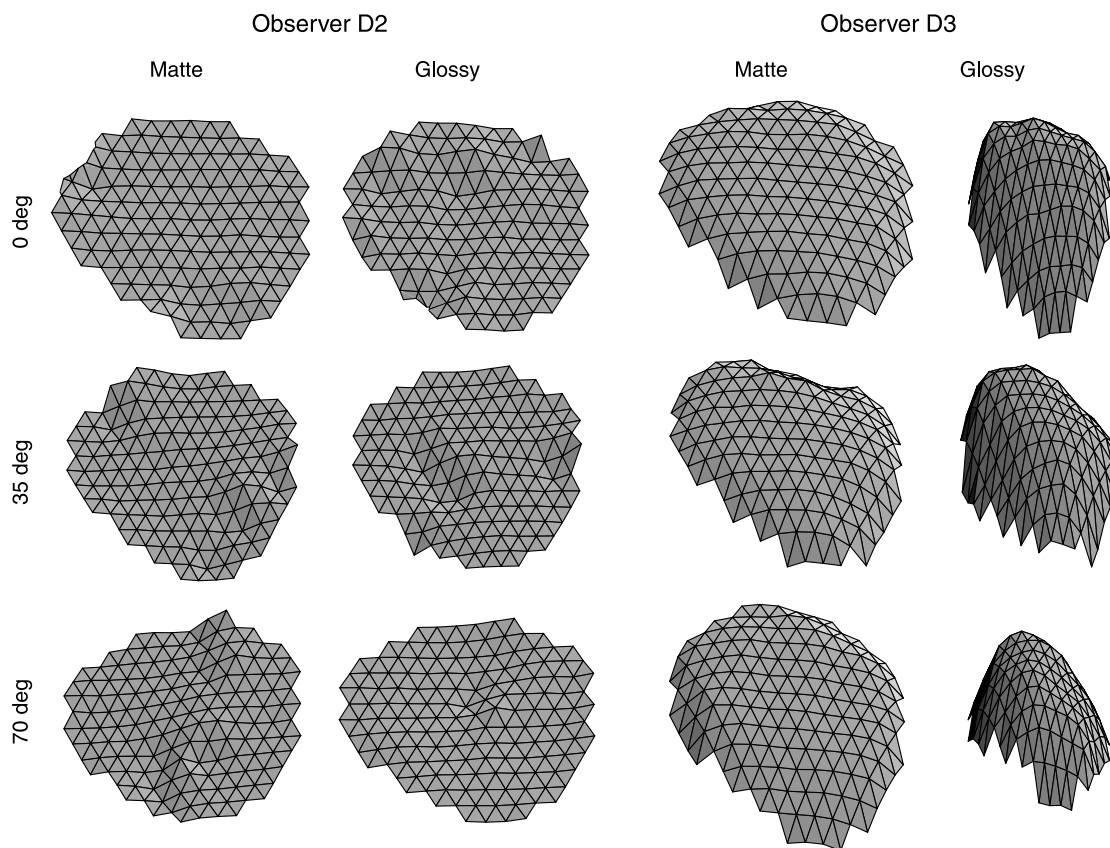


Figure 9. Results for [Experiment 4](#). Surfaces on the left are constructed from the settings of observer D2 and are representative for observers D1 and D4. Reconstructions for the settings of observer D3 are shown on the right.

markedly elongated in depth. Surface reconstructions for observer D3 are shown in [Figure 9B](#). The reconstructed surface with the most depth (1.21 screen units) in the first session in the “Glossy/70 degrees” condition and the minimum depth range (0.14 screen units) was found in session four for the “Matte/70 degrees” condition.

Discussion

The main finding in this paper is that the disparities of specular highlights and occluding contours are important variables for three-dimensional shape perception, that lead to categorically different perceived surfaces, when presented in isolation or in combination with each other. I found that contour disparity was sufficient to generate a three-dimensional percept of a connected surface as was also demonstrated earlier by Norman and Raines (2002). The ambiguities that were generated in the reduced-cue images with regard to three-dimensional shape of the depicted object are fully exploited: In some cases, the reconstructed surfaces were positively cylindrically curved and in other cases, the reconstructed surfaces were

negatively cylindrically curved. I showed in [Figure 3](#) that both reconstructions are congruent with the contour disparity and predicted that both these shapes might be observed. Specular highlights in combination with zero contour disparity on the other hand did not lead to a percept of curved surface in three out of four observers. One observer did perceive strongly curved surfaces in [Experiment 4](#). Because of the different findings in [Experiments 2](#) and [4](#), I conclude that contour disparity has strong effects on the perceived 3D shape. Disregarding the two orthogonal shape percepts in [Experiment 2](#), these results are in agreement with Norman and Raines (2002), who found that contour disparity leads to more consistent settings. The complexity of the monocularly seen contour in the objects used by Norman and Raines however already restricted the class of possible surface solutions such that two categorically different percepts, as in the present study, were not possible.

Specular highlights in combination with non-zero contour disparity led to a situation where the three specular highlights on the surface were perceived as three bumps on the surface. Perception in those cases follows the conventional interpretation of disparity rather than taking the physics of specularities and occluding contours into account. However, when presented in combination

with Lambertian shading, the disparities of specular highlights and the occluding contour seem to be ignored. Specular highlights were always in congruence with the shape specified by Lambertian shading. There was thus no cue conflict between Lambertian shading and specular highlights and we cannot be certain that specular highlights do not play a role when they were to be put in conflict. However, if specular highlights contributed to the percept in [Experiment 1](#), than they must have done differently than in [Experiment 2](#), since there was no hint of three separate bumps in any of the glossy conditions in [Experiment 1](#). That is, the interpretation of the disparity of the specular highlights depends on the presence of Lambertian shading. If there were a contribution of specular highlights, this finding would therefore suggest an interaction between the Lambertian shading cue and the specular highlights cue rather than a linear weighing of cues.

Overall, I saw a good similarity between the surface solutions from different sessions as indicated by the high R^2 's. The non-linear differences between observers as reflected in the lower R^2 's are concerned with a shift in the perceived position of the trough and peaks on the object. There is some scaling between matte and glossy conditions. The size of which is not surprising given the differences between surface solutions from different sessions. Apparently, in these data it does not matter much for the perceived depth of the surface whether a specular highlight is present or not. There is somewhat more scaling between different rotation conditions. However, the latter scaling factors also show a larger variance and lower R^2 's between observers because of which the scaling factors do not reach significance in most cases. The same can be said for shearing parameters. There is also virtually no shearing between matte and glossy conditions and somewhat more between different rotations. The latter shearing factors are again associated with higher variances between observers. Although I observed some scaling and shearing effects, these effects are small compared to the variance between sessions. There is also no obvious interaction between rotation angle and the BRDF. The position in depth of the specular highlight has apparently no major effect in this data set.

In most regressions between different rotations, the affine regression model is significantly better than the Euclidean regression model as indicated by the significant improvements of R^2 ; R^2 's did not increase significantly for regressions between matte and glossy conditions. Note that in those cases R^2 's are already in the .90's for the straight model. The lower R^2 's between different rotations are not surprising. I have observed similar non-linear behavior in earlier studies (Nefs et al., 2005).

In an earlier study on the monocular effects of specular highlights, I also found that specular highlights do not significantly increase the perceived depth over a range of

six different objects (Nefs et al., 2006). Here I extend this with the finding that the position in depth of the specular highlights does not matter. Of course we tested only for situations where the specular highlights were congruent with the same physical surface solution as the Lambertian shading. This would therefore suggest that observers either ignore the disparity of the specular highlights or interpret the disparities of specular highlights on Lambertian shading correctly and thus take the orientation of the saddle shaped surface into account. I argued before that the specific shape of the objects that were used might be an important factor for specular highlights to have an effect or not. Specular highlights tend to cling to areas of high curvature and smear in the direction of low curvature. In our experiments, the curvatures are of a gentle nature and specular highlights do not trace rims or smear across large areas of low curvature. The range of curvatures might cause the discrepancy between our results and some other studies (e.g., Todd et al., 1997). Other studies (e.g., Norman et al., 2004) also used stimuli with lower specular exponents. This causes the specular highlights to fill larger areas on the object and thus possibly reveal more information about the shape by tracing ridges, parabolic curves (e.g., Nefs et al., 2005) or smear across areas of low curvature. Similar to the monocular case where the shape of the highlight might be important, the difference between the shapes of the highlights seen by two eyes might be an important cue rather than the difference in position of the highlight on the object in the two eyes. In this study, however, I used a high specular exponent and a relatively simple object in order to favor the effects of differences in position on the object of the specular highlights in the two eyes. Under the present conditions, the specular highlights are similar in shape in the two eyes and their positions in the two eyes are clearly different. In this way I avoided the possibility that effects of the difference in position of the highlights were confounded by effects of highlight shape differences and the possibility that the disparity in the highlight position between the two eyes perhaps could not be picked up when the highlights shapes were too different from each other.

That contour disparities can be used to construct surfaces is clearly demonstrated in [Experiments 2](#) and [4](#). The stimuli give a compelling sense of a surface that is filled in from the contour inwards. Contour disparities leave the surface interpretation open for ambiguity and the visual system can both interpret the stimuli as positively or negatively curved cylindrical shape. One peculiar finding is that the reconstructed surfaces do not seem to curve to occlude themselves from sight at the edges. Rather, the surface appeared to be a cylindrical patch that is oriented in a direction that maximized absolute curvature in one direction. In [Experiment 3](#), observers were not bothered by the zero-contour disparity. That is, the monocular shape of the contour in conjunction with

shading was sufficient to lead to a stable perception of a curved surface. There was a high similarity between all reconstructed surfaces in this experiment.

Some non-linear transformations between the surface reconstruction that are observed in [Experiments 1 and 3](#) and which seem to be related to the rotation angle: Peaks and valleys seemed to be displaced depending on rotation angle. I do not believe that they are associated with contour disparity, but rather with the rotation of the shading pattern instead, since the contour disparity was zero always in any orientation in [Experiment 3](#). This is accordance with the non-linear effects we found earlier for changes in the luminance pattern and can be captured as “dark means deep” even though luminance is physically more associated with surface attitude rather than with position (Nefs et al., 2006).

Most images, including stereoscopic images, are ambiguous in their interpretation of the objects’ shapes that are depicted. These ambiguities in shading, disparity and contour have been recognized before (e.g., Anderson & Nakayama, 1994), but the implications for three-dimensional shape perception have not received that much appreciation. I demonstrated here that disparity of shading, specular highlights and the occluding contour do not on their own resolve ambiguities in stereoscopic viewing, neither mathematically nor perceptually. Their presence does not increase the perceived depth nor does it constrain the affine freedom in the surface interpretations. Lambertian shading in combination with a prior for convexity is the dominant factors in stereoscopic viewing.

Appendix A

The Method

Here I show an algorithm² to create a surface from a set of depth gradients at specified positions in the image. There are several possible variations on this method, but the principle is the same in all of them. Let us have a set of $\{x, y\}$ coordinates for the probe positions in a regular triangular tiling pattern as in [Figure A1](#). The set of coordinates is called the set of *vertices*. Each triangle of three neighboring vertices is called a *face*, and a straight connection between two neighboring vertices in a face is called an *edge*. I use a triangular tiling pattern here, but quads, hexagons or even irregular Voronoi tiling patterns would work just as well. I assigned a sequential number to each vertex and created a list of edges with pairs of vertices numbers. I also know the surface depth gradients $\{\delta z/\delta x, \delta z/\delta y\}$ at all vertices (because I have measured them) and put them in the same order as the vertices list. The depth gradients used here and the depth map that is calculated from these gradients are shown in [Figure A1](#).

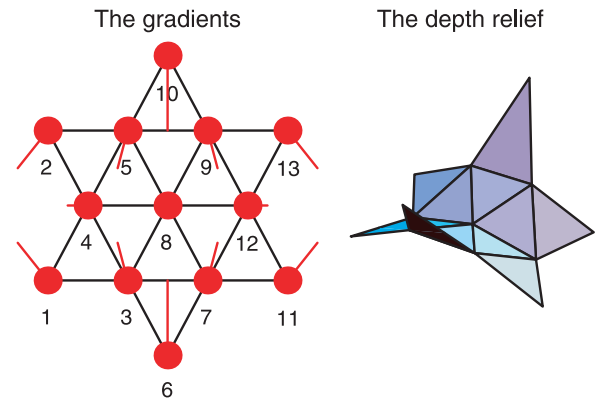


Figure A1. On the left, twelve numbered vertices with their gradients shown in red lines are shown. The length of the gradient indicates the steepness of the slope away from the observer; the direction indicates the direction of the gradient. On the right, the surface reconstruction for the gradient set on the left is shown.

Step 1

Construct an address matrix \mathbf{A} to subtract the (as yet unknown) depth values, for the two vertices of an edge. \mathbf{A} has to address all edges once:

$$\mathbf{A} \cdot \begin{bmatrix} z_1 \\ z_2 \\ z_3 \\ \dots \\ z_n \end{bmatrix} = \begin{bmatrix} z_1 - z_2 \\ z_2 - z_3 \\ z_3 - z_1 \\ \dots \\ etc \end{bmatrix}$$

$$\mathbf{A} = \begin{bmatrix} 1 & -1 & 0 & 0 & \dots \\ 0 & 1 & -1 & 0 & \dots \\ -1 & 0 & 1 & 0 & \dots \\ etc & & & & \end{bmatrix}. \quad (\text{A1})$$

Step 2

Calculate the depth difference between the two vertices (i and j) of an edge. I denote the depth difference as $\delta z_{i,j}$. Start by calculating the mean of the gradients (γ_1 and γ_2) at the two vertices. Then, take the dot product of the mean gradient, and the distance in the x - y plane between the two vertices as in [Equation A2](#). I do this for all edges and indicate this matrix as \mathbf{G} .

$$\delta z_{i,j} = \frac{\gamma_i + \gamma_j}{2} \cdot \begin{bmatrix} x_i - x_j \\ y_i - y_j \end{bmatrix}. \quad (\text{A2})$$

Step 3

Make a system of equations as in [Equation A3](#):

$$\mathbf{A} \cdot \mathbf{Z} = \mathbf{G}. \quad (\text{A3})$$

Step 4

Add one additional row to **A**, namely one that is filled entirely of 1's: At the other side of the equation I add a zero. This last addition ensures that the sum of all depth values of the surface is zero.

Step 5

Extract the array of depth values by pre-multiplying with the inverse matrix of **A**. Since **A** is not a square matrix we cannot take the inverse, but must take the pseudo-inverse of **A**. The pseudo-inverse solves the set of equations according to a least-squares principle. It is worth mentioning that this method always comes up with a solution even if you put in nonsense gradients.

$$\mathbf{A}^{-1} \cdot \mathbf{A} \cdot \mathbf{Z} = \mathbf{Z} = \mathbf{A}^{-1} \cdot \mathbf{G}. \quad (\text{A4})$$

Numeric example

A numerical example is available in the online materials that accompany this paper. I have put them in two popular mathematical packages, namely as a Mathematica notebook (numericalExample.nb), and as a MathLab m-file (numericalExample.m).

Acknowledgments

A large part of this research was supported financially by a grant from the European Union in the Information Sciences and Technologies program, contract number IST-2001-29688, a.k.a. InSight2+, while the author was at Universiteit Utrecht. An early version of the data was presented at the European Conference for Visual Perception (EVP 2005). I express thanks to Vit Drga and Julie Harris for comments on draft versions of this paper and to Jan Koenderink for advise and inspiration on this project. Also, I express many thanks to Floor van de Pavert for doing pilot experimentation.

Commercial relationships: none.

Corresponding author: Dr. Harold T. Nefs.

Email: harold.nefs@st-andrews.ac.uk.

Address: The School of Psychology, South Street, St Andrews, Fife KY16 9JP, Scotland, UK.

Footnotes

¹Shading implies here all variations in luminance depending on viewing and illumination directions. It thus

includes both Lambertian as well as non-Lambertian reflectance.

²I like to stress that I did not invent this method. Similar methods have been used much earlier (e.g., Koenderink et al., 1992) but the algorithm has never been published in a clear manner in the psychophysical literature before. I present it here merely as a tutorial service for the interested reader.

References

- Anderson, B. L., & Nakayama, K. (1994). Toward a general theory of stereopsis: Binocular matching, occluding contours, and fusion. *Psychological Review*, *101*, 414–445. [PubMed]
- Arndt, P. A., Mallot, H. A., & Bühlhoff, H. H. (1995). Human stereovision without localized image features. *Biological Cybernetics*, *72*, 279–293. [PubMed]
- Blake, A., & Bühlhoff, H. H. (1990). Does the brain know the physics of specular reflection? *Nature*, *343*, 165–168. [PubMed]
- Blake, A., & Bühlhoff, H. H. (1991). Shape from specularities: Computation and psychophysics. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *331*, 237–252. [PubMed]
- Bühlhoff, H. H. (1991). Shape from X: Psychophysics and computation. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 305–330). Cambridge, MA: The MIT Press.
- Bühlhoff, H. H., & Mallot, H. A. (1988). Integration of depth modules: Stereo and shading. *Journal of the Optical Society of America A, Optics and Image Science*, *5*, 1749–1758. [PubMed]
- Doorschot, P. C., Kappers, A. M., & Koenderink, J. J. (2001). The combined influence of binocular disparity and shading on pictorial shape. *Perception & Psychophysics*, *63*, 1038–1047. [PubMed]
- Erens, R. G., Kappers, A. M., & Koenderink, J. J. (1993). Perception of local shape from shading. *Perception & Psychophysics*, *54*, 145–156. [PubMed]
- Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision*, *4*(9):10, 798–820, <http://journalofvision.org/4/9/10/>, doi:10.1167/4.9.10. [PubMed] [Article]
- Hill, H., & Bruce, V. (1993). Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion. *Perception*, *22*, 887–897. [PubMed]
- Hill, H., & Bruce, V. (1994). A comparison between the hollow-face and 'hollow potato' illusions. *Perception*, *23*, 1335–1337. [PubMed]

- Ishihara, S. (1989). *Ishihara's tests for colour blindness, 38 plate edition*. Tokyo: Kanehara & Co, LTD.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1992). Surface perception in pictures. *Perception & Psychophysics*, *52*, 487–496. [[PubMed](#)]
- Longuet-Higgins, M. S. (1960). Reflection and refraction at a random moving surface. I. Pattern and paths of specular points. *Journal of the Optical Society of America*, *50*, 838–844.
- Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, *53*, 137–151. [[PubMed](#)]
- Nefs, H. T., Koenderink, J. J., & Kappers, A. M. (2005). The influence of illumination direction on the pictorial reliefs of Lambertian surfaces. *Perception*, *34*, 275–287. [[PubMed](#)]
- Nefs, H. T., Koenderink, J. J., & Kappers, A. M. (2006). Shape-from-shading for matte and glossy objects. *Acta Psychologica*, *121*, 297–316. [[PubMed](#)]
- Norman, J. F., & Raines, S. R. (2002). The perception and discrimination of local 3-D structure from deforming and disparate boundary contours. *Perception & Psychophysics*, *64*, 1145–1159. [[PubMed](#)]
- Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science*, *15*, 565–570. [[PubMed](#)]
- Norman, J. F., Todd, J. T., & Phillips, F. (1995). The perception of surface orientation from multiple sources of optical information. *Perception & Psychophysics*, *57*, 629–636. [[PubMed](#)]
- Oren, M., & Nayar, S. (1996). A theory of specular surface geometry. *International Journal of Computer Vision*, *24*, 105–124.
- TNO test for stereoscopic vision (8th ed.). (1972). Utrecht, The Netherlands: Laméris Instrumenten BV.
- 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*, 583–595. [[PubMed](#)]
- Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. (1997). Effects of texture, illumination and surface reflectance on stereoscopic shape perception. *Perception*, *26*, 807–822. [[PubMed](#)]
- Tse, P. U. (2002). A contour propagation approach to surface filling-in and volume formation. *Psychological Review*, *109*, 91–115. [[PubMed](#)]
- Vuong, Q. C., Domini, F., & Caudek, C. (2006). Disparity and shading cues cooperate for surface interpolation. *Perception*, *35*, 145–155. [[PubMed](#)]
- Wheatstone, C. (1852). The Bakerian lecture: Contributions to the physiology of vision. Part the second. On some remarkable, and hitherto unobserved, phenomena of binocular vision (continued). *Philosophical Transactions of the Royal Society of London*, *142*, 1–17.
- Zisserman, A., Giblin, P., & Blake, A. (1989). The information available to a moving observer from specularities. *Image and Vision Computing*, *7*, 38–42.