i-PERCEPTION

a Pion publication

i-Perception (2014) volume 5, pages 497–514

dx.doi.org/10.1068/i0645 ISSN 2041-6695 perceptionweb.com/i-perception

Is the perception of 3D shape from shading based on assumed reflectance and illumination?

James T. Todd*
Department of Psychology, The Ohio State University, Columbus, OH; e-mail: todd.44@osu.edu

Eric J. L. Egan
Department of Psychology, The Ohio State University, Columbus, OH; e-mail: egan.51@osu.edu

Flip Phillips
Psychology & Neuroscience, Skidmore College, Saratoga Springs, NY; e-mail: flip@skidmore.edu

Received 22 January 2014, in revised form 12 August 2014; published online 18 September 2014

Abstract. The research described in the present article was designed to compare three types of image shading: one generated with a Lambertian BRDF and homogeneous illumination such that image intensity was determined entirely by local surface orientation irrespective of position; one that was textured with a linear intensity gradient, such that image intensity was determined entirely by local surface position irrespective of orientation; and another that was generated with a Lambertian BRDF and inhomogeneous illumination such that image intensity was influenced by both position and orientation. A gauge figure adjustment task was used to measure observers' perceptions of local surface orientation on the depicted surfaces, and the probe points included 60 pairs of regions that both had the same orientation. The results show clearly that observers' perceptions of these three types of stimuli were remarkably similar, and that probe regions with similar apparent orientations could have large differences in image intensity. This latter finding is incompatible with any process for computing shape from shading that assumes any plausible reflectance function combined with any possible homogeneous illumination.

Keywords: 3D shape perception, shading, surface material properties.

1 Introduction

One of the most difficult problems in the study of human perception involves the ability of observers to correctly interpret patterns of image shading. The light that reflects from a visible surface toward the point of observation is influenced by three factors: (1) the surface geometry, (2) the pattern of illumination, and (3) the manner in which the surface material interacts with light. The problem for perceptual theory is to explain how it is possible to tease apart these separate influences in order to make judgments about three-dimensional (3D) shape, the pattern of illumination, or an object's material properties.

In most natural situations, the light that reflects from a local surface region toward the point of observation (i.e. luminance) varies systematically with surface orientation, as is typically described using the bi-directional reflectance distribution function (BRDF; Nicodemus, Richmond, Hsia, Ginsberg, & Limperis, 1977). Pont and Koenderink (2007) have proposed four theoretical BRDFs that represent generic types of surface materials that occur in the natural environment. These include diffuse (Lambertian) reflection on matte surfaces, specular reflection on glossy surfaces, backscattering on rough surfaces, and asperity scattering on surfaces with fine hairs such as peach skin or velvet. For a given homogeneous illumination and a given viewing direction, the BRDF describes a specific mapping between local surface orientation and luminance. This is in general a many-to-one mapping, so that it is possible for local regions with different 3D orientations to have the same luminance. However, under homogeneous illumination, all points with the same local 3D orientation must always have the same luminance.

The traditional approach to computing shape from shading is based on an assumption that all variations in luminance are due entirely to variations in surface orientation. It is typically assumed, for example, that a surface has a Lambertian reflectance function that scatters light equally in all directions, and that it is illuminated homogeneously by a collimated light field. Most models have

*Corresponding author.

assumed that the illumination direction is known. However, Kunsberg and Zucker (2013) have recently demonstrated that it is possible to obtain shape estimates that are light source invariant by analyzing the second-order differential structure of the luminance field. Although this is an important advance, their analysis still retains the traditional assumptions about Lambertian reflectance and homogeneous illumination.

There is considerable evidence to suggest that the perception of shape from shading by human observers cannot be based on a default Lambertian reflectance assumption. It was demonstrated early on that the accuracy of observers' judgments about shape or lightness are as good or better when the patterns of shading have both Lambertian and specular components (Mingolla & Todd, 1986; Norman, Todd, & Orban, 2004; Todd & Mingolla, 1983, Todd, Norman, & Mingolla, 2004). Even when the simulated displays are purely Lambertian, observers' shape judgments can deviate significantly from the expected results based on traditional models (Khang, Koenderink, & Kappers, 2007; Seyama & Sato, 1998). Of course these results do not preclude the possibility that some other assumed reflectance function exists that could account for perceptual performance.

It is possible that observers adopt a context-dependent strategy in which they first attempt to identify the surface material and then use the appropriate reflectance function to estimate 3D shape. Alternatively, they could adopt a single, hybrid reflectance function that is applied over multiple surface materials (e.g. see Wijntjes, Doerschner, Kucukoglu, & Pont, 2012). In the present article, we will present evidence that is inconsistent with either of these possibilities. In particular, we will introduce a technique called ramp shading to create stimuli for which observers' perceptions are incompatible with any plausible reflectance function combined with any possible homogeneous illumination.

A ramp-shaded image is created by texturing a surface with a planar projection of a linear intensity gradient. This defines the local image intensity for each local surface region based on its position irrespective of its orientation. Note that this is the opposite of an image rendered with a BRDF, which defines the local image intensity for each local surface region based on its orientation irrespective of its position. For surfaces that are approximately spherical, ramp-shaded images can be quite similar to those rendered with a Lambertian BRDF. This is demonstrated in Figure 1, which shows two images of a deformed sphere. The image in Panel A was rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction. The one in Panel B was textured with a linear intensity gradient that was oriented in depth (i.e. it is the equivalent of a range image). Panels C and D

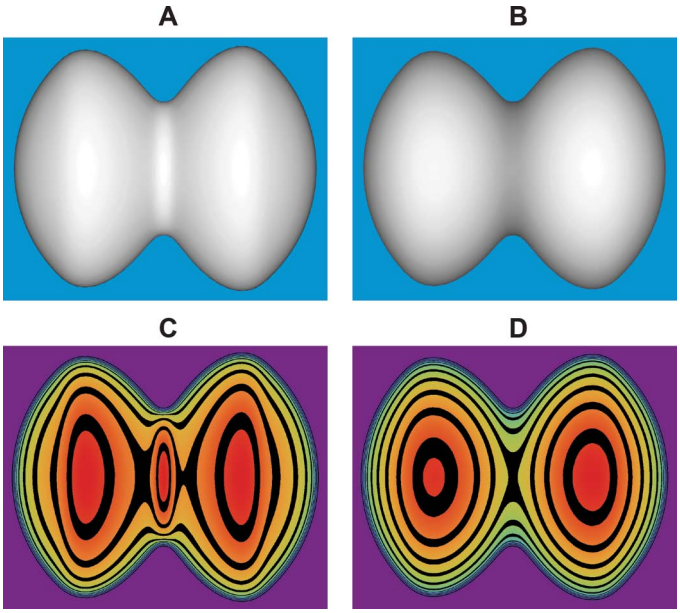

Figure 1. Two images of a deformed sphere (A and B) and their corresponding patterns of isophotes (C and D). The image in Panel A was rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction. The one in Panel B was textured with a linear intensity gradient that was oriented parallel to the viewing direction. The isointensity contour plots are coded so that red bands represent the highest image intensities and the violet bands represent the lowest.

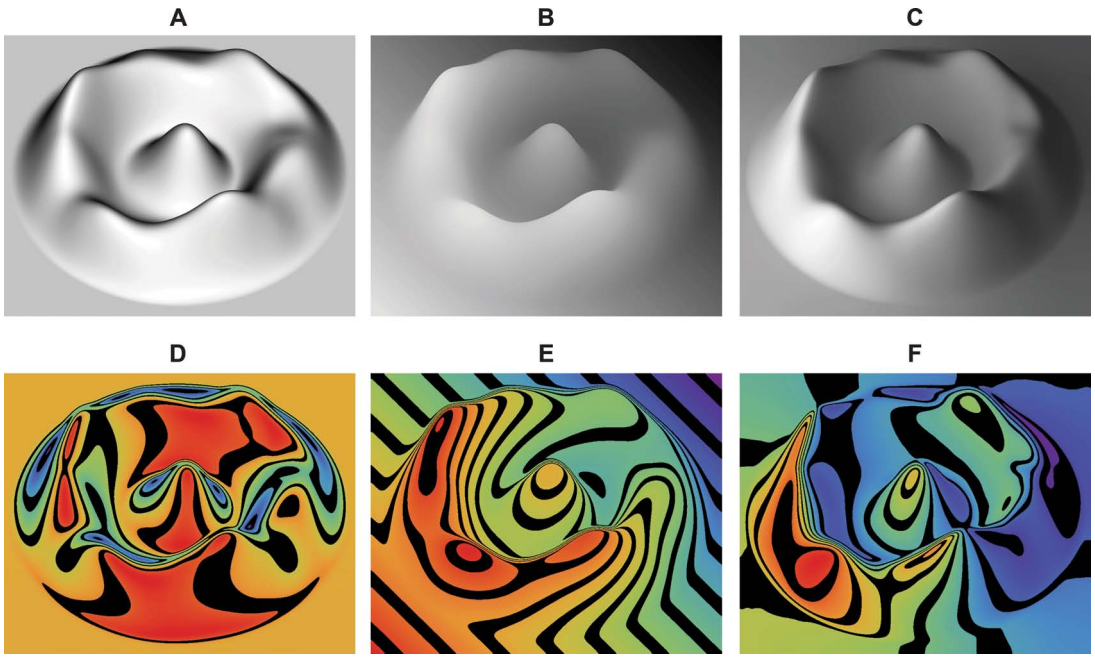

Figure 2. Three images of a deformed plane (A, B, and C) and their corresponding patterns of isophotes (C, D, and E). The image in Panel A was rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction; the one in Panel B was textured with a linear intensity gradient that was oriented at a 45° angle to each of the primary axes; the one in Panel C was rendered with a Lambertian BRDF using a large rectangular area light positioned near the lower left of the surface with two bounces of interreflection. The isointensity contour plots are color-coded so that red bands represent the highest image intensities and the violet bands represent the lowest. Observers typically report that the depicted 3D shapes in Panels A, B, and C appear quite similar, though not identical. The peaks and troughs of the ridges and valleys appear slightly less curved in the ramp-shaded image relative to the Lambertian one with homogeneous illumination, and slightly more curved in the Lambertian image with inhomogeneous illumination.

show the pattern of isophotes for these images, in which the borders between the colored bands mark points with the same image intensity, and the hue of these bands identifies the magnitudes of image intensity. Note that both images have quite similar patterns of isophotes. It is mainly in the concave crease where they diverge, such that the Lambertian shading in this region is much brighter than in the ramp-shaded version.

Lambertian and ramp-shaded images deviate more substantially when they depict distorted planes, such as the one shown in Figure 2. The image in Panel A was again rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction; the one in Panel B was textured with a linear intensity gradient that was oriented at a 45° angle to each of the primary axes; and Panels D and E show the patterns of isophotes for these images. Although ramp shading is a completely artificial technique for depicting surfaces, it is similar in some respects to translucency. The apparent glow on a translucent surface due to sub-surface scattering (e.g. on a burning candle) is attenuated by the distance that light rays must travel through the volume of the material, much like ramp shading is attenuated with distance along the gradient of the texture map. Indeed, observers often comment that the image in Panel B appears to depict a translucent material like frosted glass or snow, but that its apparent 3D shape is quite similar to the surface in Panel A.

There are other natural processes that also cause shading to vary as a function of local surface position in addition to local surface orientation. These include light attenuation with distance from the light source, cast shadows, and surface interreflections, all of which are demonstrated in Figure 2C. The surface depicted in this image has a Lambertian BRDF, and was illuminated by a large rectangular area light positioned near the lower left of the surface with two bounces of interreflection. Traditional computational models would produce distorted shape estimates for this scene because it violates the assumption of homogeneous illumination. This occurs because the ridge casts a shadow penumbra onto the valley—what is sometimes referred to as vignetting. Another thing to note in this image is the intensity gradient on the background surface due to light attenuation from the lower left to the upper

right. This would be interpreted as a curved surface by traditional models, but it is perceived correctly by human observers as an inhomogeneity in the pattern of illumination (see also Koenderink, Pont, van Doorn, Kappers, & Todd, 2007).

To better appreciate the theoretical significance of the images presented in Figures 1B, 2B, and 2C, it is useful to consider the expected outcome if observers' perceptions of shape from shading were based on an assumed BRDF and an assumed pattern of homogeneous illumination. A strong prediction of that hypothesis is that all image regions with the same apparent 3D orientation should have the same image intensity. This can be observed most clearly by examining a horizontal cross-section through the center of Figure 1A. Note that there are three points along this cross-section that appear to have a fronto-parallel orientation, and that they all have the same intensity. Contrast that with the perceptual appearance of Figures 1B, 2B, and 2C. In Figure 1B, there are also three points that have an apparent fronto-parallel orientation, but the one in the center is much darker than the other two. In Figures 2B and 2C, there are many pairs of points on opposite sides of the circular ridge that have different image intensities with the same apparent 3D orientation. This is simply not possible based on any process that computes 3D shape using an assumed BRDF with an assumed homogeneous illumination.

The research described in the present article was designed to confirm these anecdotal observations with more rigorous psychophysical measures. Observers made local orientation judgments at numerous probe points on the three shaded images shown in Figure 2. A key aspect of the experimental design is that the probe points included numerous matched pairs that both had the same depicted local orientation. For the Lambertian surface with homogeneous illumination depicted in Figure 2A, the points with the same local orientation also had the same image intensity. However, for the ramp shaded image in Figure 2B and the Lambertian surface with inhomogeneous illumination in Figure 2C, the intensities of the matched points were generally quite different.

2 Experiment 1

2.1 Methods

2.1.1 Participants

Five observers participated in the experiment, including both authors, and three others who were naïve about the issues being investigated. All of the observers had normal or corrected-to-normal visual acuity, and they all wore an eye patch to eliminate conflicting flatness cues from binocular vision.

2.1.2 Apparatus

The experiment was conducted using a Dell Dimension 8300 PC with an ATI Radeon 9800 PRO graphics card and a 19-inch gamma corrected cathode ray tube (CRT) with a spatial resolution of 1,600 × 1,200 pixels. The stimulus images were presented within a 32.5 × 32.5 cm region (1,024 × 1,024 pixels) of the display screen, which subtended 18.5° × 18.5° of visual angle when viewed at a distance of 100 cm.

2.1.3 Stimuli

The stimuli consisted of the three images shown in Figure 2. The surface depicted in Panel A was rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction; the one in Panel B was textured with a linear intensity gradient that was oriented at a 45° angle to each of the primary axes; and the surface in Panel C was illuminated by a large rectangular area light positioned to the lower left with two bounces of interreflection.

2.1.4 Procedure

The task on each trial was to adjust the slant and tilt of a circular gauge figure centered at a given probe point so that it appeared to be within the tangent plane of the surface at that point (see Koenderink, van Doorn, & Kappers, 1992, 1995). Slant is defined in this context as the angle between the surface normal and the line of sight, whereas tilt is the direction of the surface depth gradient within the fronto-parallel plane. The gauge figure simulated a small circle in 3D space with a radius of 18 pixels, and a perpendicular line at its center with a length of 18 pixels. These appeared in the image as a red ellipse with a small line along the minor axis, whose lengths and orientations could be manipulated using a handheld mouse. When adjusted appropriately, all of the observers were able to perceive this configuration as a circle oriented in depth in the tangent plane with a line perpendicular to it in the direction of the surface normal. Observers report that when performing this task they do not make quantitative

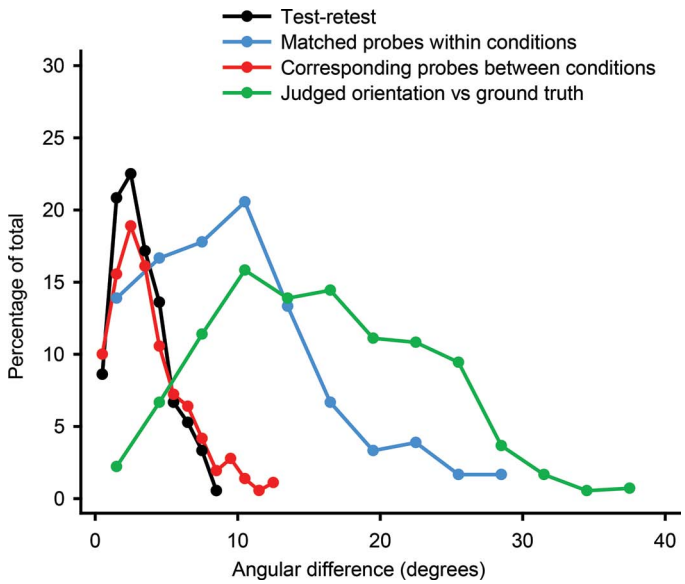

Figure 3. The distribution of test-retest differences between the first and second halves of Experiment 1 (black); the distribution of differences between corresponding probe points in all possible pairs of conditions (red); the distribution of differences between the ground truth and the average of observers' settings for each probe point in each condition (green); and the distribution of differences between the average settings of the matched probe points within a given condition (blue).

estimates of local orientation. Rather, they adjust the gauge figure so that it appears to fit on the surface, and they perform these judgments quite rapidly at a rate of 8 to 10 per minute. Another important aspect of this procedure is that it is largely unaffected by an observer's knowledge about the stimuli. For example, there have been several studies in which observers compared stimuli that had identical occlusion boundaries, but with different material properties, or presented under different viewing conditions (e.g. monocular vs. stereoscopic). These studies have revealed large differences in apparent 3D structure, even though some or all of the observers were cognitively aware that the surfaces they were comparing had the same ground truth (e.g. Khang et al., 2007; Todd, Koenderink, van Doorn, & Kappers, 1996; Todd, Norman, Koenderink, & Kappers, 1997; Wijntjes et al., 2012). Moreover, observers in these studies who do not have knowledge of the ground truth produce judgments that are comparable to those that do.

The probe points on each image included 60 pairs of regions that both had the same orientation. To select the probe points, all possible pairs of pixels were analyzed to find those whose surface normals were within 1° of one another; whose slants were between 0° and 75°; and whose image locations were separated by at least 100 pixels. Points on the planar portion of the surface were excluded from this search. The matched point pairs were sorted into five bins that spanned the range of possible slants in 15° intervals. Twelve pairs were chosen from each bin to distribute them as evenly as possible on the surface.

An experimental session was divided into three blocks, in which observers made settings for all of the different probe points within one of the possible stimulus images. Within a session, the blocks for each stimulus were presented in a random order. Each observer participated in four experimental sessions on separate days. Thus, each possible probe point in each condition was judged four times by each observer.

2.2 Results

We began our analysis by measuring the consistency of observers' judgments over multiple experimental sessions. In order to assess their test-retest reliability, we averaged the judgments over all five observers, and calculated the angular difference between their settings in the first and second halves of the experiment for each probe region in each condition. The black curve in Figure 3 shows the distribution of these test-retest differences. Note that they are all tightly clustered around a mean of only 3.2°. This provides a good estimate of the measurement error in this experiment for evaluating other findings. The red curve in Figure 3 shows the distribution of differences between the average settings at

corresponding probe points in each possible pair of conditions. Because this distribution is remarkably similar to the test-retest differences, these results indicate that any variations in the apparent shapes produced by the three types of shading were quite minimal.

The green curve in Figure 3 shows the distribution of differences between the ground truth and the average of observers' settings for each probe point in each condition. Note that the observers made large systematic errors, such that the average difference from the ground truth was 15.8°—almost five times larger than the measurement error in this experiment. Finally, the blue curve in Figure 3 shows the distribution of differences between the average settings of the matched probe points within a given condition. It is clear from these results that there were reliable differences in apparent surface orientation for many of the matched probe pairs. This has some interesting theoretical implications, because it provides strong evidence that distortions of perceived shape relative to the ground truth could not have been based entirely on the bas-relief ambiguity (Belhumeur, Kriegman, & Yuille, 1999; Koenderink, van Doorn, Kappers, & Todd, 2001). Because this ambiguity does not affect affine structure, the relative orientations of parallel regions would be invariant over all possible 3D interpretations. Thus, the fact that such regions can appear perceptually to have different orientations indicates that the relationship between apparent shape and the ground truth must include a non-affine component.

In an effort to understand the specific nature of these perceptual errors we performed affine transformations on the ground truth using the following equation:

$$1) \quad Z' = aX + bY + cZ$$

to estimate the values of the coefficients (a, b, c) that minimize the differences between the normals of the transformed surface and the observers' judgments. If observers' judgments of these surfaces had been veridical, then the values of the coefficients would be (0, 0, 1). Significant deviations from zero for the first two coefficients indicate that the apparent 3D structure is sheared relative to the ground truth. The third coefficient indicates the slope of apparent depth relative to the ground truth. For example, a slope of 0.9 would reveal that the depth relief of the surfaces was systematically underestimated by 10%. The results revealed that the best fitting values of these parameters were (0.02, −0.13, 0.59) for the Lambertian condition with inhomogeneous illumination, (−0.01, −0.13, 0.53) for the Lambertian condition with homogeneous illumination, and (0.02, −0.16, 0.53) for the ramp shading condition. We also performed a similar analysis on the average data collapsed over conditions, and the best fitting parameter values in that case were (0.02, −0.15, 0.56). These results indicate that the apparent shapes of the depicted objects were compressed in depth by roughly 44% relative to the ground truth, and that there was also a significant shear in a vertical direction. There were some variations among the different observers in this analysis, which are described in Table 1. These results are consistent with many previous studies of 3D shape from shading (e.g. Koenderink et al., 2001; Wijntjes et al., 2012) that have investigated the bas relief ambiguity. It is also important to note, however, that the average residual error with respect to the optimally transformed surface was 7.28°, which is still over twice as large as the estimated measurement error for this experiment. This provides further evidence that there was a non-affine component of the apparent distortions of the surface relative to the ground truth.

Figure 4 shows all of the probe points employed in this study superimposed on the image of the Lambertian surface with homogeneous illumination. These points are color-coded to identify the relative magnitudes of the residual errors in quintiles from the combined analysis that collapsed over conditions. The color coding of these points is based on a ROYGB convention such that points with the highest residuals are marked red, and those with the lowest are marked blue. Perhaps not surprisingly

Table 1. The best fitting coefficients (a, b, c) and residual error for the average judgments of individual observers collapsed over conditions.

Subject	Mean Residual Angle	Horizontal Shear (a)	Vertical Shear (b)	Depth Scaling (c)
EE	7.49°	–4	–13	0.50
JT	7.89°	2	–7	0.50
BL	10.47°	5	–1	0.62
CK	9.21°	–4	–16	0.59
CD	9.16°	5	–22	0.65

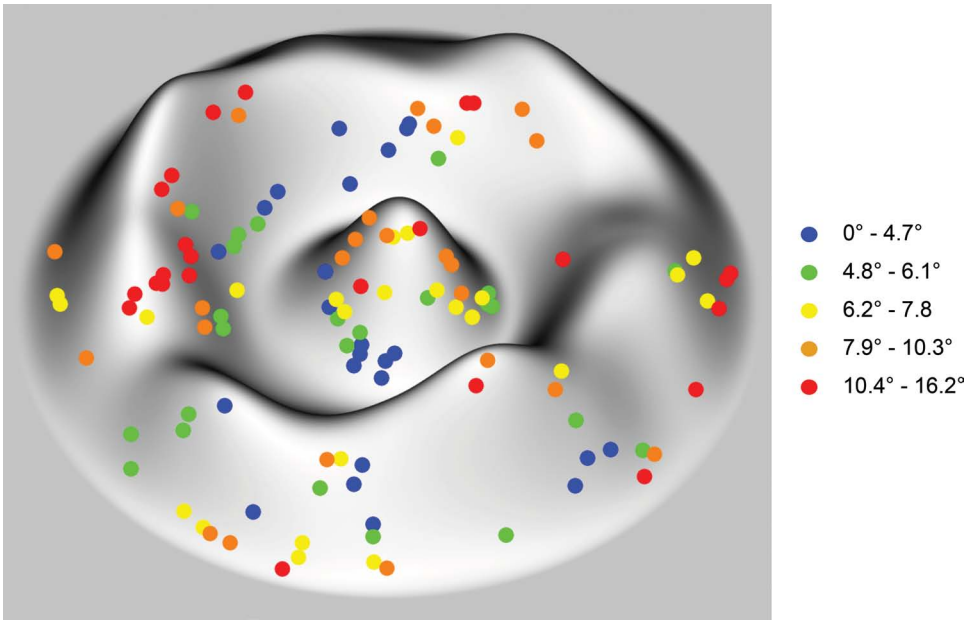

Figure 4. The residuals of the affine deformation analysis in quintiles for all of the possible probe points in Experiment 1, superimposed on the image of the Lambertian surface with homogeneous illumination.

these residuals were highly correlated with the apparent differences between the matched probe pairs ($r = 0.69$)

To what extent can traditional models of shape from shading account for these results? Let us first consider a Lambertian surface with a collimated light field that is parallel to the line of sight (see Figure 2A). Suppose that the surface has a reflectance (R), an illumination (L), and a fixed ambient component (A). For this particular special case, the image intensity (I) for any local surface region with a slant (σ) is determined by the following equation:

$$2) \quad I = RL\cos(\sigma) + A$$

After rearranging terms, this can be converted to the following:

$$3) \quad \cos(\sigma) = (I - A)/RL,$$

which provides a simple method for determining local slant from shading based on known or estimated values of R, L, and A. Note that for all possible values of these parameters, $\cos(\sigma)$ will vary linearly with I. However, our empirical results demonstrate quite clearly that there is a strong curvilinear relationship between image intensity and the cosine of observers' slant judgments (see Figure 5). These results are similar to those reported previously by Seyama and Sato (1998) and Khang et al. (2007), and they provide additional evidence that observers' shape judgments can be incompatible with an assumed Lambertian BRDF, even when the stimulus displays are consistent with that assumption.

The results presented thus far provide strong evidence that observers' judgments of local surface orientation could not have been based on an assumed Lambertian reflectance function with an assumed collimated light field that is parallel to the line of sight, but they do not eliminate the possibility that these judgments were based on some other assumed BRDF with some other pattern of homogeneous illumination. One of the primary goals of this experiment was to provide an empirical test of that possibility. If the perception of local surface orientation from shading is computed using an assumed BRDF and an assumed homogeneous illumination, then all image regions with the same apparent 3D orientation should also have the same image intensity. The present experiment was designed specifically to test that prediction by including matched probe regions whose local orientations were within 1° of one another. It is important to keep in mind, however, that our proposed test requires that the probe regions to be compared must have the same perceived orientation, as opposed to the same ground truth.

Our analysis of this issue is complicated by the fact that there were reliable differences in apparent surface orientation for many of the matched probe pairs. Fortunately, however, there were still a

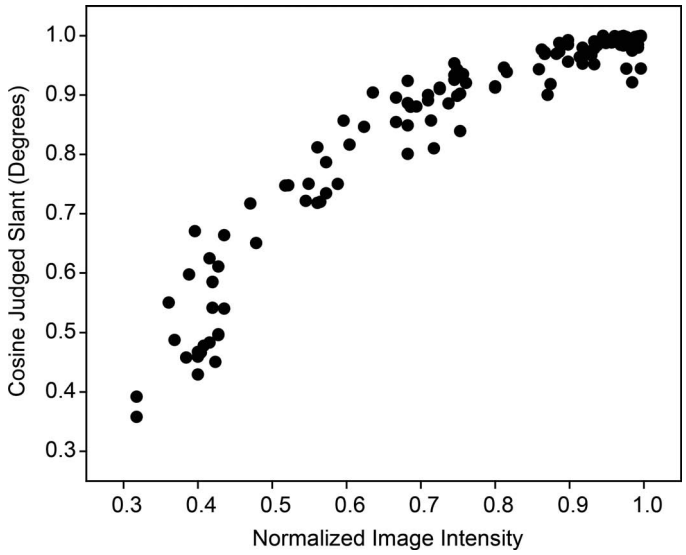

Figure 5. The cosine of adjusted slant in the homogeneous Lambertian condition of Experiment 1 as a function of image intensity.

substantial number of point pairs for which the average settings were within measurement error of one another, and we therefore restricted subsequent analyses to those for which the apparent difference in orientation was less than 6°. Almost 1/3 of the probe pairs satisfied this criterion, including 17 in the homogeneous Lambertian condition, 18 in the inhomogeneous Lambertian condition, and 21 in the ramp shading condition. The average difference in judged orientation for these restricted probe pairs was 3.3°, which is close to the average of 3.2° for the test-retest differences.

Two different analyses were performed on these restricted probe pairs. First we simulated what the differences in luminance would be for their apparent orientations using the four generic BRDFs described by Khang et al. (2007) using the following equations:

4) Diffuse scattering $I_d = (\mathbf{i}\cdot\mathbf{n})$,

5) Specular scattering $I_s = \frac{(2+k)(4+k)}{8\pi(k+2^{-k/2})}(\mathbf{h}\cdot\mathbf{n})^k(\mathbf{i}\cdot\mathbf{n})$,

6) Backscattering $I_b = \frac{(1+\mathbf{i}\cdot\mathbf{j})^k}{2^k\pi(\mathbf{i}+\mathbf{j})\cdot\mathbf{n}}(\mathbf{i}\cdot\mathbf{n})$,

7) Asperity scattering $I_a = \frac{a}{\pi[a+(\mathbf{i}\cdot\mathbf{n})(\mathbf{j}\cdot\mathbf{n})]}(\mathbf{i}\cdot\mathbf{n})$,

where **i** is the angle of incidence, **j** is the angle of exit, and **n** is the surface normal. These simulations used a random sample of 25 illumination directions with a range of slants between −40° and 40° relative to the viewing direction; and the values of the parameters a and k were varied over a range between 1 and 21 in increments of 5. For the backscattering and asperity BRDFs, the average differences in luminance were both less than 0.01 on a normalized scale from zero to one. For the diffuse BRDF, the average difference was 0.02. The highest variations of luminance occurred for the specular BRDF with an exponent of 21, which was the largest one we simulated. The average difference in that case was 0.04. Using that as a frame of reference, we then compared the differences in image intensity between the restricted probe pairs on the inhomogeneous and ramp-shaded images used in the present experiment. The average differences in these conditions were 0.17 and 0.19, respectively, which is 4.5 times larger than the simulated differences for a specular BRDF. Although it would be mathematically possible to construct an artificial BRDF that would produce large differences in luminance for small differences in surface orientation, it would be highly implausible to suggest that such a BRDF is used

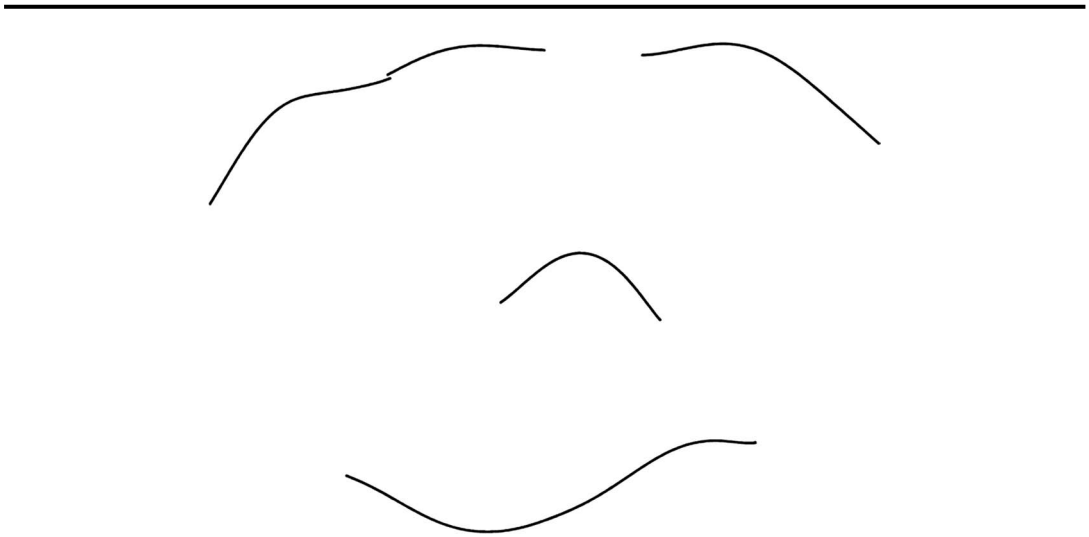

Figure 6. A toon-rendered image that depicts the occlusion contours from the stimuli in Experiment 1. This image has been cropped relative to the one that was used in the Experiment.

for the computation of 3D shape from shading in human perception. Thus, excluding that possibility, the present results provide strong evidence that observers' perceptions of these displays could not have been based on any plausible BRDF combined with any possible homogeneous illumination.

3 Experiment 2

A reviewer of an earlier draft of this article argued that the results of Experiment 1 do not necessarily indicate that observers perceive 3D shape from shading in the conditions with inhomogeneous illumination or ramp shading. According to this argument, observers may have inferred that all three of the depicted surfaces must have the same ground truth because they had identical occlusion contours. They could then have used their memories of the judged orientations in the Lambertian condition with homogeneous illumination to make their settings in the other two conditions, using the occlusion contours as an anchor to identify corresponding positions among the different stimuli. Experiment 2 was designed to test the feasibility of this strategy.

3.1 Methods

There were only two possible stimuli in this experiment: the Lambertian surface with homogeneous illumination used in Experiment 1 (see Figure 2A), and a toon-rendered version of the same surface that only depicted the occlusion boundaries (see Figure 6). These displays were judged by two of the naïve observers who participated in Experiment 1. In all other respects, the methods were identical to those reported for the previous study.

3.2 Results

When we showed the first observer the occlusion-only condition and explained the task, his immediate response was "Are you kidding?" This provided the first indication of how the results would turn out. For the judgments of the shaded image, the average test-retest difference for the two observers was 5.6° between the first and second halves of the experiment. That increased by almost three folds to 15.6° for the occlusion-only condition. Similarly, the average difference between the corresponding probe points in the different conditions was 12°, which is three times larger than the comparable differences obtained in Experiment 1. These findings show clearly that the results obtained in the inhomogeneous conditions of Experiment 1 could not have been determined based solely on the occlusion contours and observers' memories of their responses in the homogeneous Lambertian condition.

4 Experiment 3

Although observers in Experiment 2 could not reliably judge local orientation from occlusion contours presented in isolation, it does not necessarily follow that the presence of smooth occlusion boundaries in Experiment 1 had no effect at all on observers' judgments. Previous research has shown that smooth occlusion contours can provide a powerful source of information for the analysis of 3D shape from

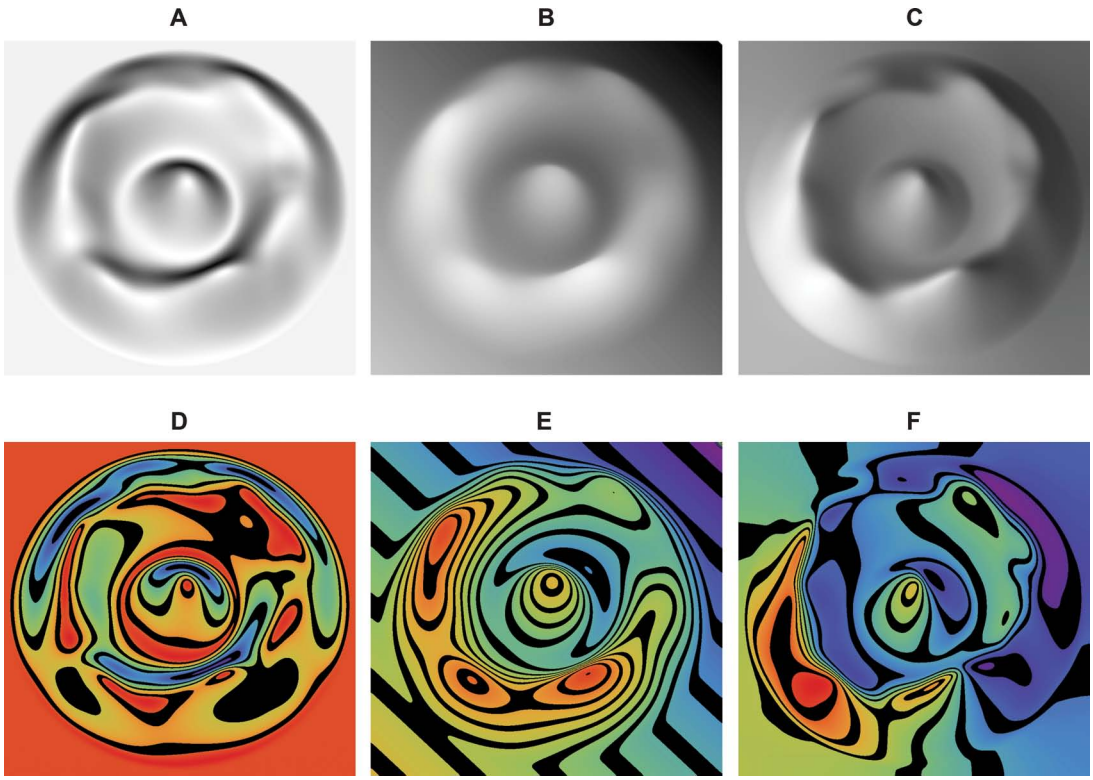

Figure 7. Three images of a deformed plane (A, B, and C) and their corresponding patterns of isophotes (C, D, and E). The image in Panel A was rendered with a Lambertian BRDF using a collimated light field that was parallel to the viewing direction; the one in Panel B was textured with a linear intensity gradient that was oriented at a 45° angle to each of the primary axes; and the one in Panel C was rendered with a Lambertian BRDF using a large rectangular area light positioned near the lower left of the surface with two bounces of interreflection. The isointensity contour plots are color-coded so that red bands represent the highest image intensities and the violet bands represent the lowest. Observers typically report that the depicted 3D shapes in Panels A, B, and C appear quite similar, though not identical. The peaks and troughs of the ridges and valleys appear slightly less curved in the ramp-shaded image relative to the Lambertian one with homogeneous illumination, and slightly more curved in the Lambertian image with inhomogeneous illumination.

shading. The surface normals along smooth occlusion contours are always perpendicular to the line of sight (Ikeuchi & Horn, 1981), which can provide a critical boundary condition for computational analyses. These contours also provide information about the surface curvature in their immediate local neighborhoods, because the sign of surface curvature in a direction perpendicular to an attached smooth occlusion contour must always be convex (Koenderink, 1984; Koenderink & van Doorn, 1982b). Under conditions of homogeneous illumination, the local luminance maxima along smooth occlusion boundaries provide additional information about the tilt of the illumination direction (Todd & Reichel, 1989).

There have been several empirical studies to show that this information can influence observers' perceptions by resolving ambiguities in the sign of surface relief (Howard, 1983; Reichel & Todd, 1990; Todd & Reichel, 1989). Consider, for example, the image presented in Figure 7A, which depicts a Lambertian surface with a collimated light field that is parallel to the line of sight. In the absence of smooth occlusion contours, the depicted relief of this surface is mathematically ambiguous. Nevertheless, for most observers, the apparent relief will be perceptually stable, because they have a strong bias to interpret the scene so that depth increases with height in the visual field. If the image is turned upside down, however, then the apparent relief can be inverted either wholly or in part (see Reichel & Todd, 1990). When surfaces are illuminated by extended light sources, there is other information from vignetting that could also be used to disambiguate the sign of surface relief (see Langer & Zucker, 1994; Langer & Bülthoff, 2000, 2001). Because concave regions see less of the light source than convex regions, the mean luminance within small concavities will be darker.

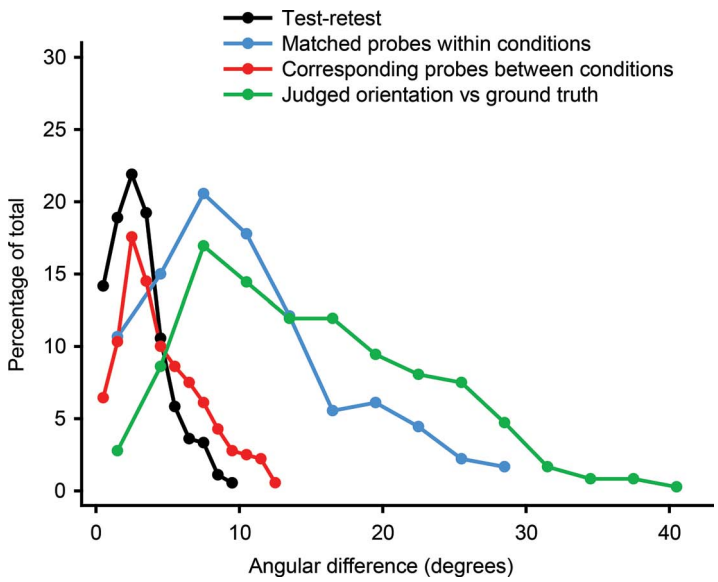

Figure 8. The distribution of test-retest differences between the first and second halves of Experiment 3 (black); the distribution of differences between corresponding probe points in all possible pairs of conditions (red); the distribution of differences between the ground truth and the average of observers' settings for each probe point in each condition (green); and the distribution of differences between the average settings of the matched probe points within a given condition (blue).

To what extent did the visible occlusion contours influence performance in Experiment 1? Experiment 3 was designed to address this issue. The stimuli were identical to those used in the earlier study, except that the overall slant of the surface was reduced just enough to eliminate all occlusions.

4.1 Methods

The methods were identical to those reported for Experiment 1 with three exceptions: (1) the three shaded images shown in Figure 7 were used as stimuli; (2) a different set of probe points were selected using the same procedure as in the previous study; and (3) these displays were judged by only four of the five observers who participated in Experiment 1. These included both authors, and two others who were naïve about the issues being investigated.

4.2 Results

The results are presented in Figure 8. The black curve shows the distribution of test-retest differences between the first and second halves of the experiment; the red curve shows the distribution of differences between corresponding probe points in all possible pairs of conditions; the green curve shows the distribution of differences between the ground truth and the average of observers' settings for each probe point in each condition; and the blue curve shows the distribution of differences between the average settings of the matched probe points within a given condition. These findings indicate that: (1) observers judgments were highly reliable; (2) variations in the apparent shapes of the depicted surfaces among the different conditions were quite minimal; (3) observers' judgments were systematically distorted relative to the ground truth; and (4) there was a non-affine component to these perceptual distortions. In other words, the results were remarkably similar to those obtained in Experiment 1. In this case, however, the observers' judgments could not have been influenced by the presence of smooth occlusion contours, because the stimuli did not contain any occlusions.

As in Experiment 1, an affine deformation analysis was performed to measure the affine components of the systematic variations between the observers' judgments and the ground truth. The best fitting parameter values for horizontal shear, vertical shear, and compression in depth were (0.01, −0.05, 0.59) for the Lambertian condition with inhomogeneous illumination, (−0.02, −0.02, 0.56) for the Lambertian condition with homogeneous illumination, and (−0.02, −0.02, 0.56) for the ramp shading condition. A similar analysis was performed on the average data collapsed over conditions, and the best fitting parameter values in that case were (0.01, −0.02, 0.56). These results are quite similar to

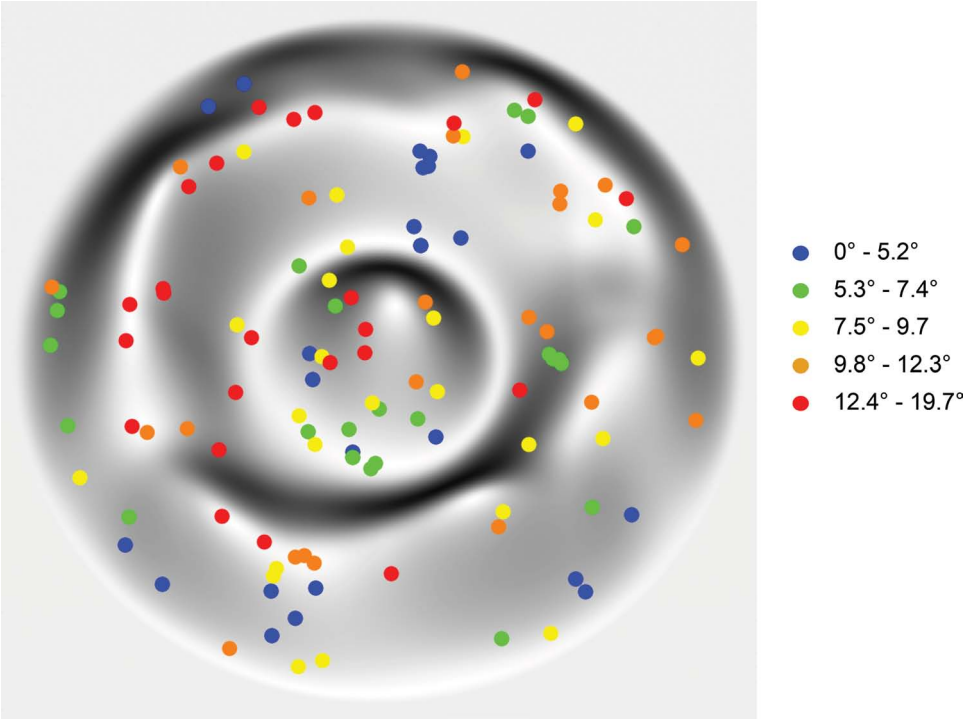

Figure 9. The residuals of the affine deformation analysis in quintiles for all of the possible probe points in Experiment 3, superimposed on the image of the Lambertian surface with homogeneous illumination.

those obtained in Experiment 1, except for a reduced amount of vertical shear. The average residual error with respect to the optimally transformed surface was 8.0°, which is over twice as large as the mean test-retest difference of 3.8° that was obtained for this experiment. This again confirms that there was a non-affine component of the apparent distortions of the surface relative to the ground truth. The magnitudes of the residuals in quintiles are shown in Figure 9 for all of the possible probe points superimposed on the image of the Lambertian surface with homogeneous illumination.

An additional analysis was performed on the restricted set of probe pairs within each condition whose judged orientations were sufficiently close to be categorized as perceptually equivalent. Because the measurement error in this study was slightly higher than in Experiment 1, we expanded the maximum difference for this set to be 7°. Over 25% of the probe pairs satisfied this criterion, including 16 in the homogeneous Lambertian condition, 19 in the inhomogeneous Lambertian condition, and 16 in the ramp shading condition. The average difference in judged orientation for these restricted probe pairs was 4.3°, which is close to the average of 3.8° for the test-retest differences. The average differences in image intensity for these restricted probe pairs in the inhomogeneous and ramp shaded conditions were 0.16 and 0.20, respectively, which is again much larger than what would be expected from any of the generic BRDFs proposed by Khang et al. (2006) with any possible homogeneous illumination.

5 Discussion

The research described in the present article was designed to compare three types of image shading: one generated with a Lambertian BRDF such that image intensity was determined entirely by local surface orientation irrespective of position; one that was textured with a linear intensity gradient, such that image intensity was determined entirely by local surface position irrespective of orientation; and another that was generated with a Lambertian BRDF and inhomogeneous illumination such that image intensity was influenced by both position and orientation. The results show clearly that observers' perceptions of these three types of stimuli are remarkably similar, even though there was little similarity in their patterns of image intensity.

It is interesting to note in this regard that there have been several other studies reported in the literature for which changes in illumination or surface material properties have not produced the same degree of shape constancy as in the present experiments (e.g. see Khang et al., 2007; Mingolla &

Todd, 1986). The primary reason for these discrepancies, we suspect, involves the depicted surface geometry. For example, in the experiments by Khang et al. (2007) and Mingolla and Todd (1986) the stimuli were restricted to ellipsoid surfaces, which appear much less perceptually compelling than the surfaces used in the present studies that contained ridges, valleys, and saddle-shaped regions. Another possible source of violations of shape constancy with changes in illumination is the use of collimated light fields. As the collimated beams become more and more slanted relative to the line of sight, the proportion of visible surface regions within attached shadows will increase. In the absence of surface interreflections, these regions will have no shading at all, and cannot therefore provide any information for the computation of 3D shape from shading (e.g. see Nefs, Koenderink, & Kappers, 2005).

A key aspect of the design of the present experiments is that the probe points included numerous matched pairs that both had the same depicted local orientation. If the perception of local surface orientation from shading is computed using an assumed BRDF and an assumed homogeneous illumination, then all image regions with the same apparent 3D orientation should also have the same image intensity. Among the sixty probe pairs we examined in each condition in each of the two experiments, approximately 30% of them had judged orientations that were within measurement error of one another. The image intensity differences between those matched probe regions in the inhomogeneous and ramp-shaded conditions were much larger than what would be expected from any of the generic BRDFs that have been proposed in the literature combined with any possible homogeneous illumination.

The present results would not necessarily be incompatible with an assumed BRDF for the computation of 3D shape from shading if the assumption of homogeneous illumination were abandoned. That latter assumption is only adopted for computational convenience, and it is almost always violated in natural vision, especially in indoor environments. However, it is not at all clear how an assumed BRDF by itself would provide sufficient constraint to compute local surface orientation from inverse optics with patterns of illumination that are inhomogeneous. Moreover, it is also difficult to reconcile that assumption with the fact that human observers can identify a wide variety of material properties. For example, consider the images presented in Figure 10, which depict a translucent milky substance, hammered gold, glossy red paint, glass, four different types of cloth and blue fur. Not only can we identify the materials in these images, but we can also perceive the depicted 3D shapes, despite the fact that they all have different inhomogeneous patterns of illumination. These observations suggest that the perception of 3D shape and material properties are somehow determined simultaneously with one another, and that they do not depend on prior knowledge about the light field.

The traditional theoretical approach to the analysis of image shading is nicely summarized in a recent paper by O'Shea, Agrawala, and Banks (2010): "Three scene properties determine the luminances in the image of a shaded object: the material reflectance, the illuminant position, and the object's shape. Because all three properties determine the image, one cannot solve for any one property without knowing the other two. Nevertheless, people perceive consistent 3D shape and consistent lighting in shaded images; they must therefore be making assumptions about the unknown properties." Computational models for determining shape from shading based on assumptions about illumination and material properties have been around since the 1970s, but their performance has been consistently disappointing (see Zhang & Tsai, 1999, for a review). The problem with all of these models is that they are designed to be used in narrowly constrained contexts, and they do not degrade gracefully when their numerous underlying assumptions are violated—as is almost always the case in natural vision.

The perceptual analysis of image shading by human observers, in contrast, is remarkably robust. Observers can correctly interpret a wide variety of optical phenomena that would wreak havoc on existing computational models, including light attenuation (e.g. Koenderink et al., 2007), cast shadows (e.g. Liu & Todd, 2004; Mamassian, Knill, & Kersten, 1998), specular highlights (e.g. Doerschner, Boyaci, & Maloney, 2010; Fleming, Dror, & Adelson, 2003; Marlow, Kim, & Anderson, 2012), transparency (e.g. Fleming, Jäkel, & Maloney, 2011), translucency (Fleming & Bülthoff, 2005), and surface interreflections (e.g. Gilchrist & Jacobsen, 1984; Madison, Thompson, Kersten, Shirley, & Smits, 2001). The present research has demonstrated, moreover, that observers can even obtain reliable information about 3D shape from images created using an artificial rendering technique that has no direct analog in natural vision. This remarkable generality is one of the most important characteristics of the ability of human observers to perceptually interpret patterns of image shading, and it should not be ignored in theoretical discussions of this phenomenon.

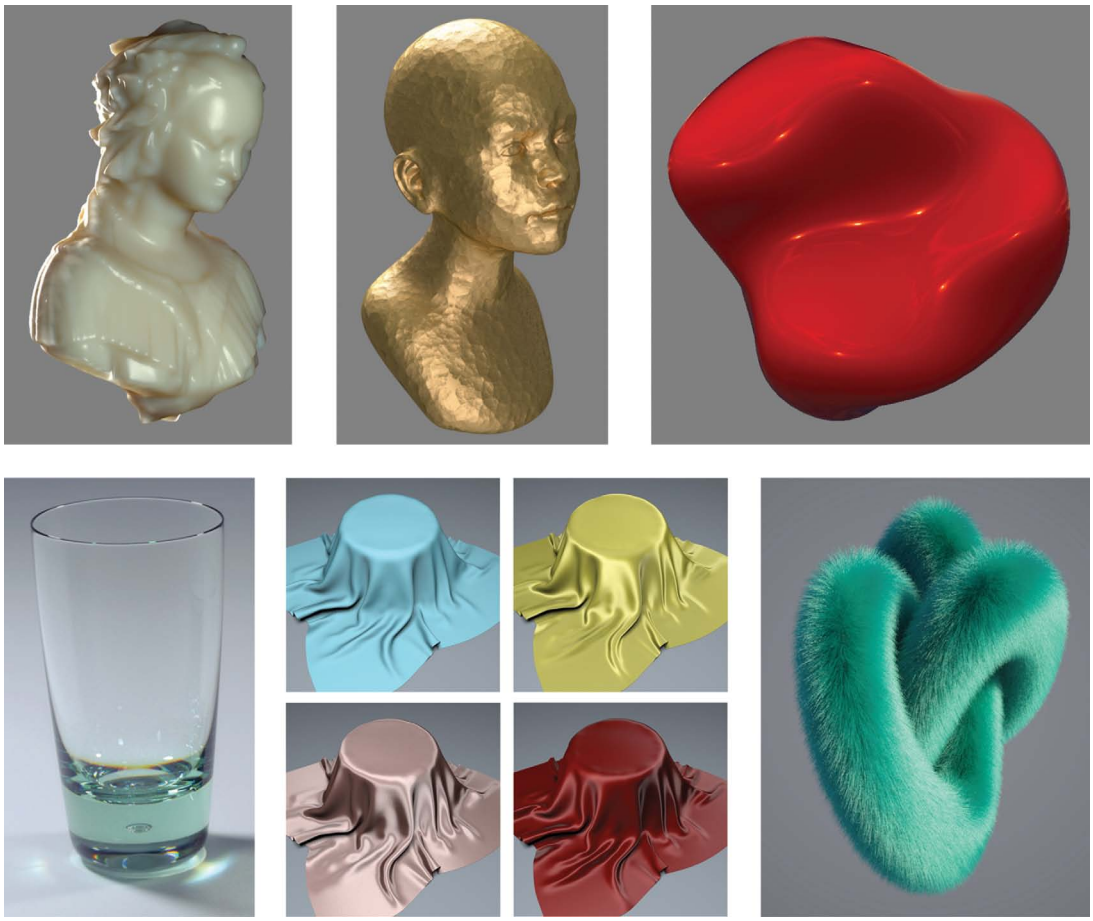

Figure 10. Example images of different 3D shapes with different material properties and different inhomogeneous illuminations. The depicted materials include a translucent milky substance, hammered gold, glossy red paint, glass, four different types of cloth, and blue fur. The glass example was created by Toni Fresnedo (tonifresnedo.com) using the Maxwell renderer. The cloth examples are from Sadeghi et al. (2013) using a new state-of-the-art algorithm for simulating cloth materials. All of the cloth examples have the same depicted geometry and the same illumination, but they have different BRDFs, which produce noticeably different patterns of shading. All of the other examples were created using the VRAY renderer. The translucent material was created using a bidirectional surface scattering reflectance distribution function (BSSRDF) for skim milk based on measurements by Jensen, Marschner, Levoy, and Hanrahan (2001).

One possible alternative to traditional models has recently been proposed by Sun and Schofield (2012). They argued that the perception of shape from shading involves two distinct modes of analysis. One they refer to as the linear shading model (Pentland, 1989) is presumed to operate for oblique illuminations such that perceived slant is proportional to local image intensity. The other they refer to as the dark-is-deep rule (Langer & Zucker, 1994). It is presumed to operate for fronto-parallel or diffuse illuminations such that perceived position in depth is proportional to local image intensity. In order to assess this hypothesis, we created images of a spherical surface illuminated by a distant point light or a hemispheric dome light at five different slants varying from 0° to 90°. We then correlated the local image intensities in these images with the local surface depths and slants on the depicted surface. The results of these analyses are presented in Table 2. Note that depth and slant both have a high negative correlation with image intensity when the primary direction of illumination is parallel to the line of sight, but that the correlations drop quite rapidly as the illumination angle is increased. Although it is conceivable that the perception of shape from shading involves multiple modes of analysis as suggested by Sun and Schofield (2012), there are very few contexts in natural vision for which a linear shading model or a dark-is-deep rule would provide an effective strategy for estimating either depth or slant (see also Langer & Zucker, 1994).

Table 2. The correlation of local image intensity with depth and slant for a spherical surface illuminated by a distant point light or a hemispheric dome light at five different slants varying from 0° to 90°.

	Correlation (r)			
	Point Light		Dome Light	
Illumination Slant (degrees)	Slant	Depth	Slant	Depth
0.00	−0.98	−1.00	−0.98	−1.00
22.5	−0.69	−0.69	−0.74	−0.75
45.0	−0.11	−0.10	−0.41	−0.43
67.5	0.34	0.32	−0.18	−0.19
90.0	0.56	0.54	0.01	0.00

Another possible alternative to traditional methods of shape reconstruction from shading was first proposed over 30 years ago by Koenderink and van Doorn (1980, 1982a). They investigated the contours that connect points of equal intensity in an image, called isophotes, in an effort to identify invariant features in these patterns that could be informative about 3D shape. Consider, for example, the X-junctions where two isophotes cross one another. For surfaces with Lambertian BRDFs and homogeneous illumination (e.g. see Figures 1A, 2A, and 7A), these X-junctions will always correspond to parabolic points on the depicted surface where there is zero curvature in one direction. A similar approach was later adopted by Breton and Zucker (1996), and, more recently, by Fleming, Torralba, and Adelson (2004). Some empirical evidence to support this approach has been reported by Wijntjes et al. (2012). They showed that two images depicting different 3D shapes appear perceptually similar if they have similar patterns of isophotes.

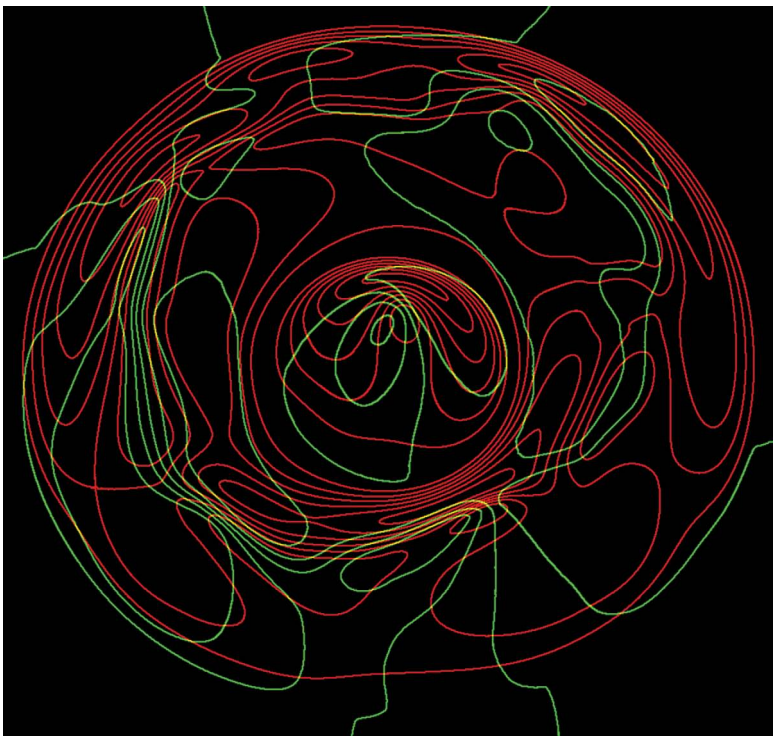

Figure 11. The isophote patterns from the homogeneous and inhomogeneous Lambertian conditions of Experiment 3, superimposed on one another in different colors. All adjacent contours in this figure represent a difference of 14.3% of the entire luminance range.

It is unlikely, however, that observers' perceptions of 3D shape from shading are based solely on the pattern of isophotes or luminance gradients, because there are some manipulations of a scene that cause large changes in the first-order image structure, but have a negligible effect on observers' shape judgments (e.g. Mingolla & Todd, 1986; Todd & Reichel, 1989). Our manipulation of the direction and manner of illumination in the present experiments provides an excellent example of this phenomenon. The effects of these lighting changes on observers' judgments were quite minimal, but they had a large effect on the first-order image structure. To demonstrate this more clearly, Figure 11 shows the isophote patterns from the homogeneous and inhomogeneous Lambertian conditions of Experiment 3, superimposed on one another in different colors. All adjacent contours in this figure represent a difference of 14.3% of the entire luminance range, such that the luminance gradients are inversely proportional to the spacing between the isophotes. Note in this case that the gradients were much steeper in the condition with a fronto-parallel collimated light field than in the one that was illuminated with a diagonally positioned area light, although the latter display produced slightly more apparent depth and curvature. Note also that in most regions of these images the isophotes (and gradients) were oriented in different directions. Based on these observations, we suspect it is the case that the higher order differential structure of an image plays a critical role in the analysis of image shading, as has also been suggested by Kunsberg and Zucker (2013). The theoretical analysis of those higher order properties remains as an important problem for future research.

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

References

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision, 35*(1), 33–44. doi:10.1023/A:1008154927611

Breton, P., & Zucker, S. (1996). Shadows and shading flow fields. *Proceedings of 1996 IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* 782–789. doi:10.1109/CVPR.1996.517161.

Doerschner, K., Boyaci, H., & Maloney, L. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. *Journal of Vision, 10*, 1–9. doi:10.1167/10.4.8.Introduction

Fleming, R., & Bülthoff, H. (2005). Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception, 2*(3), 346–382. doi:10.1145/1077399.1077409.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision, 3*(5), 347–368. doi:10.1167/3.5.3.

Fleming, R. W., Jäkel, F., & Maloney, L. T. (2011). Visual perception of thick transparent materials. *Psychological Science, 22*(6), 812–820. doi:10.1177/0956797611408734

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision, 4*(9), 798–820. doi:10.1167/4.9.10.

Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception, 13*(1), 5–19. doi:10.1068/p130005

Howard, I. P. (1983). Occluding edges in apparent reversal of convexity and concavity. *Perception, 12*(1), 85–86. doi:10.1068/p120085.

Ikeuchi, K., & Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence, 17*(1–3), 141–184. doi:10.1016/0004-3702(81)90023-0

Jensen, H. W., Marschner, S. R., Levoy, M., & Hanrahan, P. (2001). A practical model for subsurface light transport. *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '01*, 511–518. doi:10.1145/383259.383319

Khang, B.-G., Koenderink, J. J., & Kappers, A. M. L. (2007). Shape from shading from images rendered with various surface types and light fields. *Perception, 36*(8), 1191–1213. doi:10.1068/p5807

Koenderink, J. (1984). What does the occluding contour tell us about solid shape? *Perception, 13,* 321–330. doi:10.1068/p130321. Retrieved from http://www.cs.rutgers.edu/~decarlo/readings/koenderink-percep84.pdf

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2001). Ambiguity and the "mental eye" in pictorial relief. *Perception, 30*(4), 431–448. doi:10.1068/p3030

Koenderink, J. J., Pont, S. C., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2007). The visual light field. *Perception, 36*(11), 1595–1610. doi:10.1068/p5672

Koenderink, J. J., & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta: International Journal of Optics, 27*(7), 981–996. doi:10.1080/713820338

Koenderink, J. J., & van Doorn, A. J. (1982a). Perception of solid shape and spatial lay-out through photometric invariants. In R. Trappl (Ed.), *Cybernetics and systems research* (pp. 943–948). Amsterdam: North-Holland.

Koenderink, J. J., & van Doorn, A. J. (1982b). The shape of smooth objects and the way contours end. *Perception, 11*(2), 129–137. doi:10.1068/p110129.

Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1992). Surface perception in pictures. *Perception & Psychophysics, 52*(5), 487–496. doi:10.3758/BF03206710

Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1995). Depth relief. *Perception, 24*(1), 115–126. doi:10.1068/p240115

Kunsberg, B., & Zucker, S. (2013). Characterizing ambiguity in light source invariant shape from shading. *arXiv:1306.5480 [cs.CV]*. doi:10.1016/0004-3702(81)90023-0

Langer, M. S., & Bülthoff, H. H. (2000). Depth discrimination from shading under diffuse lighting. *Perception, 29*(6), 649–660. doi:10.1068/p3060

Langer, M. S., & Bülthoff, H. H. (2001). A prior for global convexity in local shape-from-shading. *Perception, 30*(4), 403–410. doi:10.1068/p3178

Langer, M. S., & Zucker, S. W. (1994). Shape-from-shading on a cloudy day. *Journal of the Optical Society of America A, 11*(2), 467–478. doi:10.1364/JOSAA.11.000467

Liu, B., & Todd, J. T. (2004). Perceptual biases in the interpretation of 3D shape from shading. *Vision Research, 44*(18), 2135–2145. doi:10.1016/j.visres.2004.03.024

Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001). Use of interreflection and shadow for surface contact. *Perception & Psychophysics, 63*(2), 187–194. doi:10.3758/BF03194461.

Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences, 2*(8), 288–295. doi:10.1016/S1364-6613(98)01204-2.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology: CB, 22*(20), 1909–1913. doi:10.1016/j.cub.2012.08.009

Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics, 53*, 137–151. doi:10.1007/BF00342882

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., & Limperis, T. (1977). Geometrical considerations and nomenclature for reflectance. *Science And Technology, 160*(October), 1–52. doi:10.1109/LPT.2009.2020494

Nefs, H. T., Koenderink, J. J., & Kappers, A. M. L. (2005). The influence of illumination direction on the pictorial reliefs of Lambertian surfaces. *Perception, 34*(3), 275–287. doi:10.1068/p5179

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*(8), 565–570. doi:10.1111/j.0956-7976.2004.00720.x

O'Shea, J., Agrawala, M., & Banks, M. (2010). The influence of shape cues on the perception of lighting direction. *Journal of Vision, 10*, 1–21. doi:10.1167/10.12.21.Introduction

Pentland, A. (1989). Shape information from shading: A theory about human perception. *Spatial Vision, 4*(2–3), 165–182. doi:10.1109/34.24791

Pont, S. C., & Koenderink, J. J. (2007). Matching illumination of solid objects. *Perception & Psychophysics, 69*(3), 459–468. doi:10.1167/10.9.5.

Reichel, F. D., & Todd, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception and Performance, 16*(3), 653–664. doi:10.1037/0096-1523.16.3.653

Sadeghi, I., Bisker, O, Decken, J. D. & Jensen, H.W. (2013). A practical microcylinder appearance model for cloth rendering. *ACM Transactions on Graphics, 32*(2):14, 1–12. doi:10.1145/2451236.2451240.

Seyama, J., & Sato, T. (1998). Shape from shading: estimation of reflectance map. *Vision Research, 38*(23), 3805–3815. doi:10.1016/S0042-6989(97)00435-5.

Sun, P., & Schofield, A. J. (2012). Two operational modes in the perception of shape from shading revealed by the effects of edge information in slant settings. *Journal of Vision, 12*, 1–21. doi:10.1167/12.1.12

Todd, J. T., Norman, J. F., & Mingolla, E. (2004). Lightness constancy in the presence of specular highlights. *Psychological Science, 15*(1), 33–39. doi:10.1111/j.0963-7214.2004.01501006.x

Todd, J. T., Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1996). Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception and Performance, 22*(3), 695–706. doi:10.1007/s00426-008-0145-7.

Todd, J. T., & Mingolla, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance, 9*(4), 583–595. doi:10.1037/0096-1523.9.4.583

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*(7), 807–822. doi:10.1068/p260807.

Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review, 96*(4), 643–657. doi:10.1037/0033-295X.96.4.643

Wijntjes, M. W. A., Doerschner, K., Kucukoglu, G., & Pont, S. C. (2012). Relative flattening between velvet and matte 3D shapes: Evidence for similar shape-from-shading computations. *Journal of Vision, 12*(1):2, 1–11. doi:10.1167/12.1.2

Zhang, R., & Tsai, P. (1999). Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21*(8), 690–706.

James Todd studied psychology and computer science at the University of Connecticut, where he received his Ph.D. in 1974. He is currently a professor at the Ohio State University in the Department of Psychology. His primary research interests include the visual perception of three-dimensional form from various types of optical information, such as shading, texture, motion, and binocular disparity, and the visual control of motor actions.

Eric Egan is a Post Doctoral Researcher under the supervision of James Todd. Eric received his B.A. in neuroscience at Skidmore College in 2008 and his Ph.D. in psychology at The Ohio State University in 2014. His primary research interests include the visual perception of three-dimensional shape from both visual and haptic information.

Flip Phillips began his career at Pixar before returning to Ohio State for a Ph.D. in cognitive psychology. There, he specialized in the perception of three-dimensional shape, inspired by his earlier architectural and computer graphics training. A member of the Skidmore faculty since 1998, his research centers on the visual and haptic perception of two- and three-dimensional shapes as well as psychological aesthetics.