

# The effects of field of view on the perception of 3D slant from texture

James T. Todd \*, Lore Thaler, Tjeerd M.H. Dijkstra

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

Received 19 October 2004; received in revised form 21 December 2004

## Abstract

Observers judged the apparent signs and magnitudes of surface slant from monocular textured images of convex or concave dihedral angles with varying fields of view between 5° and 60°. The results revealed that increasing the field of view or the regularity of the surface texture produced large increases in the magnitude of the perceptual gain (i.e., the judged slant divided by the ground truth). Additional regression analyses also revealed that observers slant judgments were highly correlated with the range of texture densities (or spatial frequencies) in each display, which accounted for 96% of the variance among the different possible dihedral angles and fields of view.

© 2005 Elsevier Ltd. All rights reserved.

## 1. Introduction

In an influential series of articles that were first published over 50 years ago, Gibson (1950a, 1950b) introduced the concept of texture gradients as a potential source of optical information for the perceptual specification of 3D surface structure. The image in Fig. 1 provides a compelling example of the perceptual salience of this information. Note how the pattern of texture elements produces the appearance of a smoothly curved surface with well defined regions of concavity and convexity. There are several local measures of optical texture that have been discussed in the literature as possible cues for the analysis of 3D shape (e.g., see Gårding, 1992), and the left panel of Fig. 2 defines four of the most common ones, which will be referred to here as length ( $\lambda$ ), width ( $\omega$ ), area ( $\alpha$ ), and foreshortening ( $\phi$ ). The right panel of Fig. 2 shows how these measures vary systematically in an image as a function of surface depth and orientation. Gibson argued that it is the pattern of these systematic variations (i.e., the gradients)

that provides useful information about 3D surface structure.

In an effort to model this intriguing aspect of human perception, researchers in computational vision have developed numerous mathematical algorithms that are designed to estimate the 3D structure of surfaces from patterns of optical texture (e.g., Aloimonos, 1988; Blake & Marinos, 1990; Brown & Shvaytser, 1990; Clerc & Mallat, 2002; Davis, Janos, & Dunn, 1983; Gårding, 1992, 1993; Kanatani & Chou, 1989; Malik & Rosenholtz, 1997; Purdy, 1958; Super & Bovik, 1995; Witkin, 1981). It is interesting to note, however, that the vast majority of these algorithms do not make use of gradient information, as originally suggested by Gibson. Rather, they estimate local slant from the foreshortening of individual texture elements or a statistical sample of surface markings within a limited local neighborhood. These models assume that variations in reflectance on a surface are statistically isotropic—i.e., that they are approximately equal in all directions. Thus, any deviation from isotropy in the projected image of a texture can be used to estimate the direction and magnitude of local surface slant.

It is only quite recently that models have been developed that come closer to Gibson's original conception

\* Corresponding author. Tel.: +1 614 292 8661; fax: +1 614 292 5601.  
E-mail address: [todd.44@osu.edu](mailto:todd.44@osu.edu) (J.T. Todd).




Fig. 1. A pattern of optical texture that is perceptually interpreted as a smoothly curved 3D surface.




Fig. 2. Local properties of optical texture. The left panel defines four different attributes of optical texture elements that could potentially provide useful information about the 3D structure of an observed surface. Variations of texture length are often referred to in the literature as scaling gradients, and, when the elements are aligned in depth, as linear perspective or gradients of convergence. The term compression has also been used by some authors to describe texture width, whereas others use this same term to describe texture foreshortening. The right panel shows how these local texture attributes are affected by changes in surface depth and orientation. Note that changes in orientation have a minimal effect on texture length, and that changes in depth have a minimal effect on texture foreshortening.

by explicitly analyzing gradients of optical texture (Clerc & Mallat, 2002; Gårding, 1992, 1993; Malik & Rosenholtz, 1997). These gradient based algorithms estimate surface slant or curvature by comparing patterns of optical texture across different local neighborhoods. For example, the model developed by Malik and Rosenholtz (1997) is based on an assumption that the texture on a physical surface is statistically homogeneous—i.e., that it is invariant over translation—and it estimates the 3D shape of a surface by measuring local affine deformations of the texture within neighboring patches of a visual image. This approach works well for planar or singly curved surfaces, but it is not easily generalized to doubly curved surfaces, except in the special case where curvature is approximately the same in all directions.

Which class of models is most representative of human perception? Fig. 3 provides a simple demonstration that may be helpful for evaluating this issue. When the image in the left panel is presented to naïve observers, they almost always report a perceptual appearance of an approximately planar surface that is slanted in depth. In the

absence of other evidence, it cannot be determined whether the apparent slant of this surface is due to the foreshortening of individual texture elements or the systematic variations in the sizes and shapes of these elements. However, it is possible to manipulate these different sources of information independently of one another. As will be described in more detail below, the variations in optical texture on planar surfaces are primarily determined by the depicted field of view (FOV). For example, the image in the left panel of Fig. 3 was rendered with a camera angle of  $60^\circ$ . The right panel of this figure shows exactly the same surface at exactly the same viewing distance. The only differences are that the depicted field of view has been decreased to  $5^\circ$ , and the size and spacing of the polka dots has been reduced so that the visible surface region has approximately the same number of elements as in the left panel.<sup>1</sup> Note that the texture elements are still noticeably foreshortened, but that the reduced field of view has effectively eliminated any detectable texture gradients. When the image in the right panel is presented to naïve observers, they never report the perceptual appearance of a surface slanted in depth, thus suggesting that the presence of visible gradients may be essential for the perception of 3D shape from texture.

To better appreciate the various factors that can influence the magnitudes of texture gradients, it is useful to consider an infinite planar surface that is covered with circular polka dots. Suppose that an observer is positioned at a vantage point  $P_0$ , which is a distance  $D$  from the closest point on the surface (see left panel of Fig. 4). If a circular texture element with a diameter  $L$  were centered on that point, its optical projection would have an arc length  $\lambda_0$  as defined by the following equation:

$$\lambda_0 = 2 \arctan(L/2D) \quad (1)$$

Let us now consider an arbitrary surface point  $P$  with an optical slant  $\sigma$ . Provided that  $\lambda_0$  is smaller than about  $30^\circ$ , the optical properties of a texture element with a diameter  $L$  that is centered on  $P$  will be closely approximated by the following equations:<sup>2</sup>

<sup>1</sup> Another way of characterizing the differences between the left and right panels of Fig. 3 is that they have different degrees of perspective. The perspective of a scene is typically defined as a ratio between the depth of its nearest point and the depth of its farthest point (Braunstein & Payne, 1969). When a planar surface is observed from a fixed vantage point, the magnitude of perspective varies systematically with the field of view.

<sup>2</sup> Eqs. (2)–(6) are based on a type of projective geometry called weak perspective, in which distortions of optical texture elements as a function of changes in depth or orientation are limited to affine transformations (see also Malik & Rosenholtz, 1997). For example, in weak perspective the optical projection of a square texture element would always be a parallelogram, whereas in strong perspective it could also be a trapezoid. In most natural viewing contexts, the differences between weak and strong perspective are negligible. They can diverge significantly, however, when individual texture elements have large angular extents.




Fig. 3. Two images of a planar surface that were both rendered from the same viewing distance with a 60° FOV (left) and a 5° FOV (right). The size and spacing of the polka dots have been adjusted so that each pattern contains approximately the same number of elements.




Fig. 4. Factors that can affect the variations of optical texture on planar surfaces. The left panel depicts a planar surface that is viewed from a vantage point  $P_0$ . A texture element with a diameter  $L$  and a viewing distance  $D$  is positioned at a point on the surface whose orientation is perpendicular to the line of sight. An arbitrary surface point  $P$  is also depicted that has an optical slant  $\sigma$ . The right panel shows the pattern of variation of  $\lambda$  and  $\phi$  for planar surfaces over all possible values of  $\sigma$ , and the range that is visible over a limited field of view.

$$\text{Foreshortening } (\phi) = \cos(\sigma) \tag{2}$$

$$\text{Length } (\lambda) = \lambda_0 \cos(\sigma) \tag{3}$$

$$\text{Width } (\omega) = \lambda_0 \cos^2(\sigma) \tag{4}$$

$$\text{Area } (\alpha) = \pi \lambda_0^2 \cos^3(\sigma) / 4 \tag{5}$$

The right panel of Fig. 4 shows the pattern of variation of  $\lambda$  and  $\phi$  for planar surfaces over all possible values of  $\sigma$ . It is of course not possible for surfaces in the natural environment to have an infinite extent, so they must always be observed with a limited field of view, such as the one that is depicted in Fig. 4 by the two dashed lines. It should be evident from this figure that for planar surfaces the range of variation in  $\lambda$  and  $\phi$  is determined by two parameters: The field of view and the optical slant  $\sigma$  at its center. Thus, from the perspective of a gradient based analysis, a reasonably accurate estimate of surface slant from texture should only be possible when these parameters are sufficiently large for the variations

in texture to be reliably detected (see also Blake, Bühlhoff, & Sheinberg, 1993).

Why would observers use texture gradients as a source of information for surface slant, rather than a zero-order property like foreshortening, which can be measured more reliably? A fundamental shortcoming of all zero-order measures of texture is that they are ambiguous with respect to sign. For example, an analysis of surface slant based solely on local foreshortening would be incapable of distinguishing whether a surface is slanted upward or downward. It is only in the spatial derivatives of optical texture that this ambiguity can be resolved. Gradient information is also essential for determining the direction and magnitude of surface curvature from texture (Gårding, 1992).

Given the potential importance of field of view for the information content of textured images (Blake et al., 1993), it is somewhat surprising that there has been so little research on how variations in viewing angle affect observers' perceptions of 3D slant from texture. We

know of only two previous experiments that have addressed this issue,<sup>3</sup> and both of them were restricted to a very limited set of conditions. For example, Knill (1998a) measured slant discrimination thresholds for planar surfaces with random ellipse textures at four possible fields of view between 3° and 25°, but all of his displays had a fixed standard slant of 65°. Similarly, Tibau, Willems, Van Den Berg, and Wagemans (2001) measured the magnitudes of apparent slant for planar surfaces with square grid textures, but they only employed a small range of viewing angles between 4.8° to 8°. Because of the limited scope of these studies, the research described in the present article was designed to provide a more systematic investigation of how the perception of shape from texture is influenced by variations in the field of view. The stimuli depicted convex or concave dihedral angles whose planar faces had mean optical slants between 25° and 65°. These surfaces were presented with several different types of texture, and a wide range of viewing angles between 5° and 60°. Observers indicated the perceived shape of each depicted surface using an adjustment task, which made it possible to measure the magnitude of perceived slant as well as its precision.

## 2. Methods

### 2.1. Subjects

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

### 2.2. Apparatus

The experiment was conducted using a Dell Dimension 8300 PC with an ATI Radeon 9800 PRO graphics card. Images that were rendered with a camera angle of 20° or less were presented on a standard CRT, and they had horizontal and vertical extents of 30.0 cm. Images that were rendered with a camera angle greater than 20° were back projected onto a translucent display

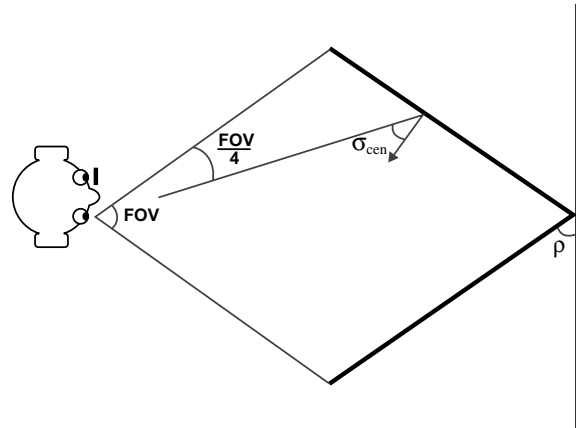


Fig. 5. A schematic representation of the basic scene geometry that was depicted in the stimuli of the present experiment.

screen using an LCD projector, and they had horizontal and vertical extents of 121.9 cm. Both types of display had a spatial resolution of 1280 × 1024 pixels. The viewing distances were adjusted appropriately over a range of possible values between 85 and 344 cm so that the visual angle of each image would be matched to the camera angle with which it was rendered. The displays were viewed monocularly with an eye patch, and a chin rest was used to constrain head movements.

### 2.3. Stimuli

A schematic representation of the basic scene geometry is presented in Fig. 5. The stimuli all depicted concave or convex dihedral angles that were bilaterally symmetrical about a vertically oriented edge. Variations in the depicted 3D structures of these displays were controlled by two parameters: The field of view (FOV), and the optical slant ( $\sigma_{\text{cen}}$ ) in the center of each face. There were five possible values of FOV (5°, 10°, 20°, 40° and 60°) and five possible values of  $\sigma_{\text{cen}}$  (25°, 35°, 45°, 55° and 65°) that were presented in all possible combinations. Because the optical slants for planar surfaces vary linearly with visual direction, with a slope of one,<sup>4</sup> the maximum and minimum values of optical slant were defined by the following equations:  $\sigma_{\text{max}} = \sigma_{\text{cen}} + \text{FOV}/4$ , and  $\sigma_{\text{min}} = \sigma_{\text{cen}} - \text{FOV}/4$ . It follows, therefore, that the range of optical slants for the concave and convex dihedral angles were perfectly matched. They were not matched, however, with respect to the physical slants

<sup>3</sup> A third set of studies by Buckley, Frisby, and Blake (1996) manipulated the field of view over a range of values between 10° and 30°. When texture cues were all consistent with one another, the manipulation of FOV had no detectable effect on observers' slant judgments. However, this result is difficult to interpret with respect to the monocular perception of slant from texture, because the texture information in all of their displays was combined with slant information from binocular disparity. Other researchers have also investigated the effects of field of view on the perception of planar surfaces from motion (Cornilleau-Pérès et al., 2002) and binocular disparity (Bradshaw, Glennerster, & Rogers, 1996; Rogers & Bradshaw, 1995).

<sup>4</sup> The proof of this is quite simple. Let the origin ( $\varpi_0$ ) for defining visual direction be the direction that is perpendicular to an observed planar surface (see Fig. 4). In that case, the relation between visual direction ( $\varpi$ ) and optical slant ( $\sigma$ ) is given by the following equation:  $\varpi = \sigma$ . The first derivative of this equation is:  $d\sigma/d\varpi = 1$ , which is invariant over all possible viewing distances and orientations. Note that the derivative is unaffected by changing the origin for defining visual direction.

( $\rho$ ) of the depicted surfaces relative to the fronto-parallel plane. These were defined by the following equations:  $\rho = \sigma_{\text{cen}} + \text{FOV}/4$  for concave surfaces, and  $\rho = \sigma_{\text{cen}} - \text{FOV}/4$  for the convex surfaces.

Images of the dihedral angles were rendered using 3D Studio Max by Kinetix with five possible textures that were created with the DarkTree 2.0 texture plugin by Darktree Studios. These textures included (1) a plaid pattern, (2) a pattern of regular contours in the direction of surface slant, (3) a pattern of irregular contours in the direction of slant whose widths were modulated by a random noise function, (4) a pattern of regular (i.e., circular) blobs that were positioned at random without overlapping, and (5) a pattern of irregular blobs that varied in size and shape. The textures were scaled for each combination of  $\sigma_{\text{cen}}$  and FOV so that they would all have approximately 10 texture elements along a vertical cross-section through the center of each planar face. Examples of the different types of stimuli are shown in Figs. 6–10. Each of these figures contains four images of a surface with a mean optical slant ( $\sigma_{\text{cen}}$ ) of  $65^\circ$  that all have the same texture. Convex and concave dihedral angles are shown in the left and right columns, respectively. Images in the top row have a depicted FOV of  $60^\circ$ , whereas those in the bottom row have a depicted FOV of  $10^\circ$ . These images are perceptually most compelling when viewed monocularly.

### 2.4. Procedure

The field of view was manipulated across blocks so that all of the trials in a given experimental session had the same viewing distance. On each trial, an image of a dihedral angle was presented on one of the two possible display screens. A second monitor was located off to the side of the main display that contained an adjustment figure, which observers could manipulate with a hand held mouse to match the apparent cross section in depth of the depicted surface. Because many of the stimuli did not appear to have sharp edges (e.g., see Figs. 8–10), the adjustment figure was a hyperbola rather than a dihedral angle in order to provide a closer match to the observers’ perceptions. The shape of the adjustment figure was defined by the following equation:

$$y^2 = b^2 + x^2 \tan^2 \rho \tag{6}$$

where  $\rho$  is the angle of the asymptotic lines relative to the horizontal direction, and  $b$  is the distance from the apex of the hyperbola to the intersection of its asymptotic lines, which is also referred to as the “semi-transverse axis”. The adjustment figure had a fixed width of 5.71 cm, and a height that varied from 0 to 13.2 cm. Observers were informed that the vertical ( $y$ ) axis of the adjustment figure was intended to represent the apparent depth of the depicted surface and they were instructed to adjust the shape of this figure so that it




Fig. 6. Examples of surfaces with plaid textures used in the present experiment. Convex surfaces are shown on the left and concave surfaces on the right. The images in the upper and lower rows have depicted fields of view of  $60^\circ$  and  $10^\circ$ , respectively.






Fig. 7. Examples of surfaces with regular contour textures used in the present experiment. Convex surfaces are shown on the left and concave surfaces on the right. The images in the upper and lower rows have depicted fields of view of  $60^\circ$  and  $10^\circ$ , respectively.




Fig. 8. Examples of surfaces with irregular contour textures used in the present experiment. Convex surfaces are shown on the left and concave surfaces on the right. The images in the upper and lower rows have depicted fields of view of  $60^\circ$  and  $10^\circ$ , respectively.




Fig. 9. Examples of surfaces with regular blob textures used in the present experiment. Convex surfaces are shown on the left and concave surfaces on the right. The images in the upper and lower rows have depicted fields of view of 60° and 10°, respectively.




Fig. 10. Examples of surfaces with irregular blob textures used in the present experiment. Convex surfaces are shown on the left and concave surfaces on the right. The images in the upper and lower rows have depicted fields of view of 60° and 10°, respectively.




Fig. 11. A schematic diagram of the adjustment screen used in the present experiment.

matched the apparent shape of a horizontal surface cross section as closely as possible.

The overall design of the adjustment screen is presented in Fig. 11. At the beginning of each trial, the observers were required to click on one of three buttons to indicate whether the surface shown in the main display appeared convex, concave or perfectly flat, and this response was used to constrain the orientation of the adjustment curve. Next they adjusted the parameter  $\rho$  with a horizontal mouse movement to indicate the apparent slants of the two planar faces of the depicted dihedral angle. Finally, they adjusted the parameter  $b$  (i.e., the semi-transverse axis) with a scroll bar to indicate the apparent curvature of the dihedral edge. Once observers were satisfied with their settings, they could move on to the next trial by clicking on a button that was labeled “next”. It was also possible to move backward in the sequence to modify a previous response, though none of the observers reported that they made use of that option. All observers agreed that these response tasks were quite natural and that they had a high degree of confidence in their judgments.

### 2.5. Design

To summarize the overall experimental design, there were 250 possible conditions: 2 signs of curvature (concave or convex)  $\times$  5 possible fields of view ( $5^\circ$ ,  $10^\circ$ ,  $20^\circ$ ,  $40^\circ$  and  $60^\circ$ )  $\times$  5 possible values of the mean optical slant  $\sigma_{\text{cen}}$  ( $25^\circ$ ,  $35^\circ$ ,  $45^\circ$ ,  $55^\circ$  and  $65^\circ$ )  $\times$  5 different textures (plaids, regular contours, irregular contours, regular blobs and irregular blobs). Within a given experimental session, the field of view remained fixed, and the 50 pos-

sible combinations of curvature, slant and texture were presented once each in a random sequence. Each observer participated in four separate sessions for each of the five possible fields of view.

## 3. Results

All of the observers reported at the conclusion of the experiment that they experienced dramatic variations in the relative magnitudes of perceived depth and slant across the different field of view conditions. With the  $5^\circ$  field of view, many of the depicted surfaces appeared perfectly flat or nearly so. With the  $40^\circ$  or  $60^\circ$  fields of view, in contrast, the observers all reported that the appearance of 3D shape was as perceptually compelling as binocular stereopsis.

### 3.1. The perceived sign of curvature

For the plaid and regular contour textures, the observers' sign of curvature judgments were 100% accurate for all of the different combinations of  $\sigma_{\text{cen}}$  and FOV. For the remaining textures, the accuracy was reduced to 94% for the irregular contours, 86% for the regular blobs and 76% for the irregular blobs. Among these latter three textures, the overall pattern of performance was quite similar in that the incorrect judgments of the depicted sign of curvature occurred primarily in those conditions with the lowest values of  $\sigma_{\text{cen}}$  and FOV. Fig. 12 shows the combined performance for the blob and irregular contour textures as a function of optical slant, for each of the five possible fields of view. Note how the manipulation of these parameters had a huge effect on the ability to distinguish concave and convex dihedral angles. For the lowest values of  $\sigma_{\text{cen}}$  and FOV the accuracy of observers' judgments was not significantly different from chance, whereas at the highest

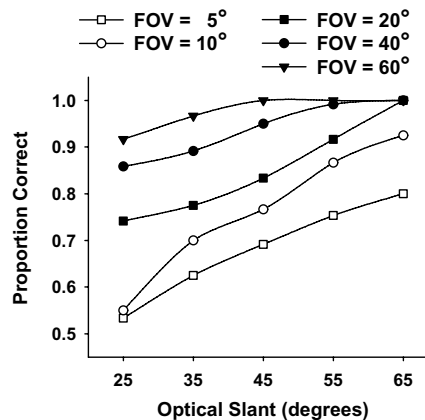


Fig. 12. The proportion of correct sign of curvature judgments for the blob and irregular contour textures as a function of optical slant for each possible field of view.



values of  $\sigma_{cen}$  and FOV, their judgments were 100% accurate.

### 3.2. Variance of the slant settings

To evaluate the precision of observers' slant settings, we calculated the standard deviations among the four repeated judgments for each observer in each of the 250 possible display conditions. Over all observers and conditions, the average value of these standard deviation measures was  $4.66^\circ$ . An analysis of variance revealed that there were no significant differences among the different conditions except for a significant effect of optical slant,  $F(4, 16) = 8.598, p < .001$ , and a significant effect of the sign of surface curvature,  $F(1, 4) = 8.948, p < .05$ , both of which are shown in Fig. 13.

### 3.3. Magnitudes of perceived slant

Fig. 14 shows the average judged physical slant over all observers and textures as a function of the ground

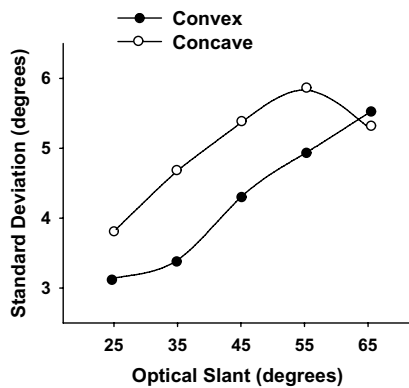


Fig. 13. The standard deviations of observer slant settings as a function of optical slant for the concave and convex dihedral angles.

truth for each possible sign of curvature and field of view. For almost all of the different possible combinations of FOV, texture type and sign of curvature, there was a strong linear correlation between the observers' settings and the depicted physical slants of the dihedral faces. The field of view also had a large effect on the magnitude of perceived slant, though there was an interesting interaction between field of view and the sign of surface curvature that we had not expected. Note in Fig. 14 that for the convex surfaces, each successive increase in FOV produced a corresponding increase in the magnitude of observers' slant settings. A similar pattern of performance was obtained for the concave surfaces up to a FOV of  $20^\circ$ , but increases in FOV beyond  $20^\circ$  had little or no effect on the magnitude of perceived slant.

In order to facilitate a statistical analysis of these data, we converted the raw physical slant judgments to a measure of perceptual gain, which is defined as the ratio of judged slant relative to the ground truth. Fig. 15 shows the average perceptual gain as a function of field of view for each possible sign of curvature and texture. One important aspect of these results that deserves to be highlighted is the systematic variation in the accuracy of observers' slant judgments for the different types of texture. The highest average perceptual gain of 0.74 was obtained for surfaces with plaid textures. For the remaining conditions, the gain was reduced to 0.67 for the regular contours, 0.56 for the regular blobs, 0.52 for the irregular contours, and 0.32 for the irregular blobs. An analysis of variance confirmed the significance of all of the different patterns of variation that are evident in Figs. 14 and 15. That is to say, the analysis revealed significant main effects of texture type,  $F(4, 16) = 9.926, p < .001$ , field of view,  $F(4, 16) = 94.95, p < .001$ , and the sign of surface curvature,  $F(1, 4) = 68.422, p < .001$ , and a significant interaction

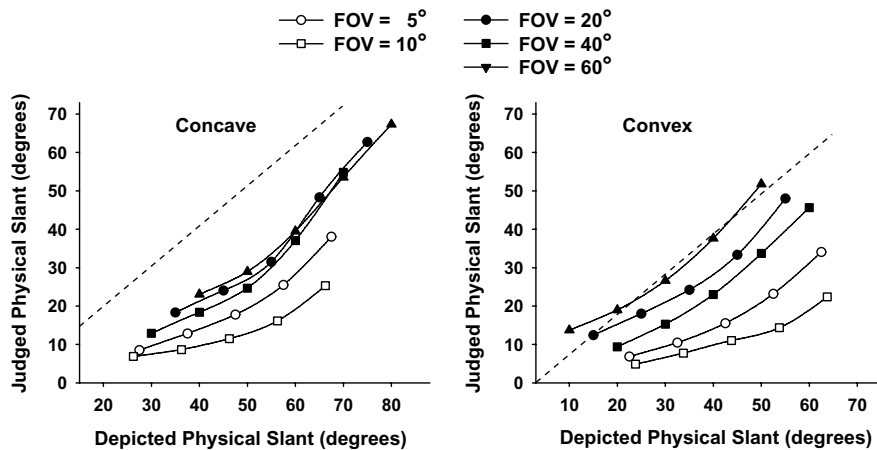


Fig. 14. The average physical slant settings as a function of depicted physical slant for each possible field of view and sign of curvature. The dashed lines represent veridical performance.

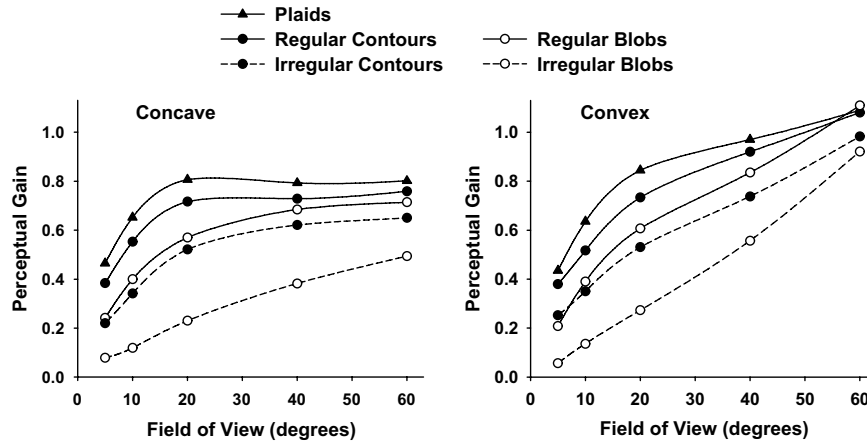


Fig. 15. The average perceptual gain as a function of field of view for each possible texture and sign of curvature.

between field of view and the sign of surface curvature  $F(4, 16) = 45.513, p < .001$ .

Additional regression analyses were performed in an effort to determine the specific attributes of texture on which observers' slant judgments were based. Preliminary plots of the data revealed that the best fits would be obtained using log-log coordinates. Thus, the analysis began by computing the logarithm of adjusted slant, averaged over observers and textures, for each of the 50 possible combinations of curvature, FOV and optical slant. These transformed slant settings were then correlated with the logarithms of numerous measures of optical texture that varied systematically across the different conditions. For example, these measures included the maximum texture length ( $\lambda_{max}$ ) within each stimulus, the minimum texture length ( $\lambda_{min}$ ), the median texture length  $(\lambda_{max} + \lambda_{min})/2$  and the range of texture lengths  $(\lambda_{max} - \lambda_{min})$ . Similar correlations with adjusted slant were also performed on the maximum values, minimum values, median values and ranges of texture width ( $\omega$ ), area ( $\alpha$ ), vertical density ( $1/\lambda$ ), horizontal density ( $1/\omega$ ) and total density ( $1/\alpha$ )—note that the density measures can also be interpreted as spatial frequencies. Texture foreshortening ( $\phi$ ) was not included in this analysis because that measure was linearly proportional to texture length. It is important to point out that these regression analyses did not make use of actual image measures. Rather, they were computed indirectly from local opti-

cal slants ( $\sigma$ ), based on the fact that  $\lambda$  and  $\phi$  are proportional to  $\cos(\sigma)$ ,  $\omega$  is proportional to  $\cos^2(\sigma)$  and  $\alpha$  is proportional to  $\cos^3(\sigma)$ —see Eqs. (2)–(5).

The results of these analyses are presented in Table 1, which shows the squared coefficient of regression ( $r^2$ ) for each of the correlations described above. Note that the best fits of the data were obtained for the range measures of density ( $1/\lambda$ ,  $1/\omega$  and  $1/\alpha$ ), all of which accounted for at least 90% of the between display variance. For purposes of comparison, the ground truth accounted for only 49% of the variance. While examining the scatter plots of these correlations, we noticed that most of the residual variance was due to lower than expected slant judgments for the convex surfaces with 40° or 60° fields of view, which may have been caused by an inability of the visual system to process high frequency information in peripheral regions of the visual field. In an effort to correct for this, we recalculated the range measures such that any information outside the central 30° of view would be ignored for the convex surfaces. This modification produced significantly improved fits for all of the density range measures. The best fit of all was obtained for the vertical density range (see Fig. 16), whose correlation with adjusted slant accounted for 96% of the variance among the 50 possible combinations of curvature, FOV and optical slant. This same measure also provided excellent fits for the individual textures, though there were systematic differences in

Table 1

The squared coefficients of regression ( $r^2$ ) between the logarithms of judged slant collapsed over observers and textures, and the logarithms of several different measures of optical texture that varied among the different stimulus conditions

	Length ( $\lambda$ )	Width ( $\omega$ )	Area ( $\alpha$ )	Vertical density ( $1/\lambda$ )	Horizontal density ( $1/\omega$ )	Total density ( $1/\alpha$ )
Minimum value	0.73	0.73	0.73	0.28	0.28	0.28
Median value	0.62	0.56	0.29	0.65	0.66	0.67
Maximum value	0.28	0.28	0.28	0.73	0.73	0.73
Range	0.61	0.34	0.12	0.90	0.91	0.90

Note that the different density measures can also be interpreted as spatial frequencies. See text for details.

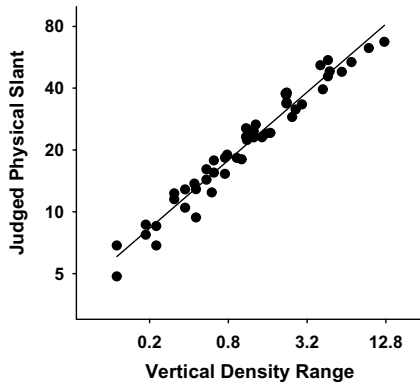


Fig. 16. The average judged slant as a function of the vertical density range. In an effort to provide an intuitively clear interpretation of this measure, the units along the horizontal axis have been calibrated to reflect the number of visible cycles for the regular contour textures shown in Fig. 7. The vertical density range in this context is the unsigned difference between the number of visible cycles at the outer edges of each display relative to the number of visible cycles in its center. For convex surfaces, the range measure was limited to the central 30° of view.

the slopes of the regression lines. The values of  $r^2$  and slope, respectively, were 0.89 and 0.41 for the plaid textures, 0.92 and 0.40 for the regular contours, 0.91 and 0.48 for the irregular contours, 0.90 and 0.74 for the regular blobs and 0.84 and 1.06 for the irregular blobs.

### 3.4. The perceived curvature of the dihedral edge

The apparent curvature ( $\kappa$ ) of the dihedral edges in each of the different conditions was computed from the observers' settings of  $\sigma$  and  $b$  using the following equation:  $\kappa = (\tan^2\sigma)/b$ . In all of the various conditions depicting surfaces with plaid textures, observers always set the semi-transverse axis ( $b$ ) of the adjustment figure to a value of zero, thus indicating that the apparent curvature of the dihedral edge had a value of infinity. This

was also mostly the case for surfaces with regular contour textures, although two of the observers produced some non-zero settings in conditions with low values  $\sigma_{cen}$  and FOV. For the blob and irregular contour textures, however, the dihedral edges were almost never perceived as discontinuities. Fig. 17 shows that averaged judged curvatures for these textures as a function of optical slant for each possible sign of curvature and field of view. These data show clearly how the judged curvatures of the dihedral edges varied with optical slant and FOV in a remarkably systematic manner. There was also a large difference between the judged curvatures of the convex and concave surfaces—i.e., note that the scales of the two graphs differ by a factor of five. As can be observed in the example stimuli presented in Figs. 8–10, the edges of the concave surfaces appear to have much higher curvatures than the edges of the convex surfaces.

## 4. Discussion

Before considering the theoretical implications of these results, it is useful to review some other relevant findings that have been reported previously in the literature. Tibau et al. (2001) were the first to report that the apparent slant of a planar surface can be significantly influenced by the depicted field of view. Their displays all had square grid textures, and the FOV ranged from 4.8° to 8°. Tibau et al.'s stimuli were most similar to the plaid texture displays in the present experiment, and the results we obtained for those textures were quite similar to theirs with comparable fields of view.

Another related experiment that deserves to be highlighted was performed by Knill (1998a), who measured slant discrimination thresholds for images of planar surfaces with random ellipse textures using varying fields of view that ranged from 3° to 25°. When the lowest value

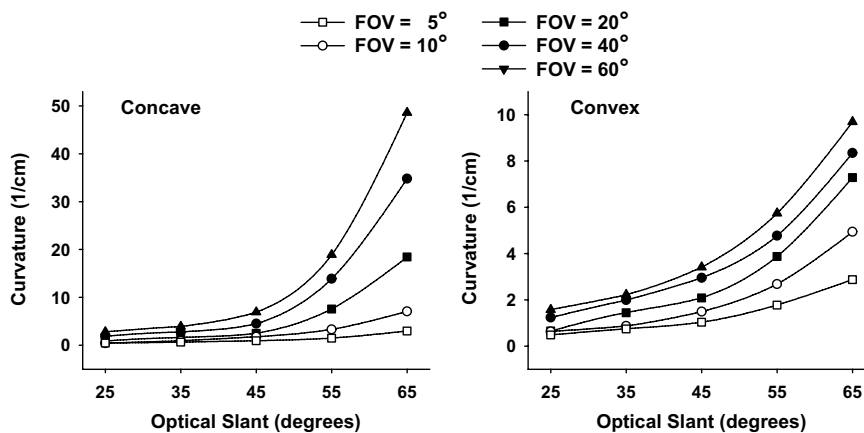


Fig. 17. The average judged curvature of the dihedral edges for the blob and irregular contour textures as a function of optical slant for each possible field of view and sign of curvature.




Fig. 18. Perspective projections of a planar surface that were rendered with a  $3^\circ$  field of view. The one on the left has a depicted slant of  $70^\circ$ , and the one on the right has a depicted slant of  $50^\circ$ . Note that both surfaces appear to have a fronto-parallel orientation, despite the fact that the texture elements are highly foreshortened.

of optical slant depicted in these images had a fixed value of  $52.5^\circ$ , reductions in the viewing angle produced an approximately twofold increase in the observer's thresholds. However, when the highest value of optical slant had a fixed value of  $77.5^\circ$ , the manipulation of FOV had no effects whatsoever—even when the viewing angle was reduced to only  $3^\circ$ . Because systematic variations of optical texture are negligible in such small field displays, this latter finding could be interpreted as strong evidence that an accurate perception of surface slant from texture need not require gradient information. There is a potential problem with that interpretation, however. Consider the images presented in Fig. 18 of planar surfaces with a  $70^\circ$  slant (left) and a  $50^\circ$  slant (right), which were both rendered with a  $3^\circ$  field of view. These images are clearly discriminable because of differences in texture foreshortening and density, yet when observers are asked to make magnitude estimations of the apparent slants, they almost always report that both surfaces appear to have a fronto-parallel orientation. If texture patterns with different degrees of foreshortening can have the same apparent slant, then the evaluation of Knill's results becomes more complicated. Although his observers were explicitly instructed to make judgments about slant, we must also consider what strategies they may have employed when the depicted slants were perceptually indistinguishable. It is of course possible that they may have guessed randomly in that situation, but we suspect it is more likely that they may have relied on predictive 2D cues, like foreshortening, whenever such cues were available.

A third line of research that is closely related to the present experiment has been performed by Li and Zaidi (2000, 2001, 2003) and Zaidi and Li (2002). They investigated the perceived signs of curvature of sinusoidally corrugated surfaces with a wide variety of different types of texture. These surfaces were presented with 1–3 visible cycles, but the individual cycles all had fields of view between  $4^\circ$  and  $7^\circ$ . The results of these studies have con-

sistently indicated that observers' judgments are almost perfectly accurate for surfaces with contour textures in the direction of maximum curvature (e.g., see Figs. 6–8), but that the overall level of performance is no better than chance for surfaces with isotropic textures (e.g., see Figs. 9,10).<sup>5</sup> On the basis of these findings, Li and Zaidi have argued that isotropic textures are inherently inadequate for accurately specifying the sign of surface curvature or the direction of surface slant. In considering this conclusion, it is important to note that the present experiment replicated their findings under comparable viewing conditions. With a  $5^\circ$  field of view, the observers were 100% accurate at judging the sign of curvature for surfaces with plaid or regular contour textures, and they were well below threshold for surfaces with isotropic blob textures. However, when we consider a broader range of viewing conditions, our results show clearly that this particular pattern of performance is limited to small fields of view. For example, in the  $20^\circ$  field of view conditions the observers were 90% accurate at judging the sign of curvature for surfaces with regular blob textures, and they were 99% accurate in the  $40^\circ$  and  $60^\circ$  FOV conditions.

Why should some textures be more difficult than others for determining the sign of surface curvature or the direction of surface slant? From the perspective of a gradient based analysis, these effects of texture type are most likely caused by the relative difficulty of distinguishing systematic variations of optical texture from random variations that occur due to noise. For example,

<sup>5</sup> Although the overall levels of performance obtained by Li and Zaidi with isotropic textures were no greater than chance, the results showed clearly that observers did not make their responses by guessing randomly about the sign of surface curvature. Rather, the pattern of errors indicates that observers have a systematic bias to interpret surfaces as convex (see also Langer & Bühlhoff, 2001; Liu & Todd, 2004), such that the curvatures of convex surfaces are consistently judged correctly, and concavities are consistently misinterpreted as convexities.

the measurement of texture length gradients is easiest for textures that contain linear contours in the direction of surface slant. Note in Figs. 6 and 7 how the distance  $\lambda$  between the horizontal contours varies as a smoothly continuous linear function, which provides noise free information about the optical length gradient. This is also referred to in the literature as linear perspective or as the gradient of convergence. The measurement of texture gradients is inherently more difficult for random blob textures, because information from spatially separated texture elements must somehow be interpolated in order to achieve the perceptual representation of a smoothly continuous surface. This is likely to be especially difficult when there are large random variations in the sizes and shapes of the texture elements, as was the case for the irregular blob textures of the present experiment (see Fig. 10). Although these random variations can be averaged out by pooling information over a sufficiently large neighborhood, this would cause the systematic variations in optical texture to become blurred, which would reduce their relative magnitudes, and, therefore, their detectability. The spatial pooling of information for the measurement of local texture properties is also most likely to be responsible for the reduced apparent curvatures of the dihedral edges that were obtained for the blob and irregular contour textures.

#### 4.1. On the relative salience of different texture cues

Much of the previous literature on the perception of 3D shape from texture has focused on whether some local attributes of texture are more perceptually informative than others (e.g., length, width, area, density or foreshortening). Most researchers who have attempted to address this issue have used some sort of cue conflict paradigm, in which different attributes of texture are independently manipulated, either within static monocular images (Attneave & Olson, 1966; Cutting & Milard, 1984; Phillips, 1970; Todd & Akerstrom, 1987), or stereoscopic displays in which the texture information can also be in conflict with binocular disparity (Buckley et al., 1996; Cumming, Johnston, & Parker, 1993; Frisby & Buckley, 1992). In general, the results of these studies have indicated that for random blob textures without linear perspective, observers rely most heavily on foreshortening information, and to a lesser extent on texture length (i.e., scaling), but that texture density has little or no influence on the perception of 3D shape (see Knill, 1998a for a review).

In light of these past results, it is especially surprising that observers' judgments in the present experiment were most highly correlated with the range measures of texture density by a large margin over other texture cues. How can this be reconciled with the previous literature? It is important to keep in mind when considering this

issue that cue conflict paradigms have been criticized because they may induce observer strategies that are not representative of normal perceptual processing (Blake et al., 1993; Knill, 1998a, 1998b), or because the texture properties that are systematically manipulated—usually length and foreshortening—are inevitably confounded with other local properties like width or area that could potentially be important for observers' judgments (Todd & Akerstrom, 1987). Manipulations of texture density are particularly problematic in this regard, because they produce systematic variations in the relative proportions light and dark pixels across different regions of an image, which would ordinarily remain constant for homogeneous textures in natural vision (see Todd & Reichel, 1990).

The results of the present experiment are perhaps less surprising if the density measures in Table 1 are reinterpreted as measures of spatial frequency, because previous studies have also suggested that systematic variations in spatial frequency may be a primary source of information for judgments of 3D shape from texture (Li & Zaidi, 2003; Todd & Akerstrom, 1987). Given what is known about the structure of the human visual system, it is highly implausible that texture patterns could be perceptually analyzed by counting the number of individual texture elements within local neighborhoods of an image or by precisely measuring their lengths, widths or aspect ratios. A more likely possibility, we suspect, is that the analysis of texture gradients is performed primarily in the Fourier domain (e.g., Malik & Rosenholtz, 1997) or by non-linear networks of spatially tuned filters (e.g., Grossberg & Mingolla, 1985, 1987). The correlation of judged slant with the systematic variation of spatial frequency as reported in Table 1 is quite compatible with either of those two approaches.

#### 4.2. Can results with planar surfaces be generalized to curved surfaces?

Because planar surfaces are ubiquitous in both natural and man-made environments, they are an obviously important special case for the perceptual analysis of 3D surface structure. There is some evidence to suggest, however, that the results obtained with planar surfaces in the perception of 3D shape from texture may not always generalize to other forms of surface geometry. For example, the mathematical analysis described in the introduction showed that systematic variations of optical texture on planar surfaces are primarily determined by the depicted field of view, but that is not the case for smoothly curved surfaces. For surfaces that are sufficiently curved, it is possible to obtain large systematic variations among optical texture elements over relatively small angular extents, and these variations can produce compelling perceptions of 3D surface structure.





Fig. 19. Images of smoothly curved surfaces under orthographic projection.

Whereas textured planar surfaces observed with small fields of view produce little or no apparent depth (see Figs. 3 and 18), the perception of 3D shape for smoothly curved surfaces can be obtained even under orthographic projection, which is mathematically equivalent to rendering a scene from an infinite viewing distance with an infinitely small camera angle. Fig. 19 shows two examples of smoothly curved textured surfaces that were each rendered under orthographic projection (see also Knill, 2001; Todd & Oomes, 2002; Todd & Reichel, 1989, 1990). Naïve observers almost always report that these images are perceived as 3D surfaces,<sup>6</sup> and previous psychophysical research has shown that the magnitude of perceived depth for this type of display is only slightly attenuated relative to the apparent depths that are produced when the same surfaces are presented under strong polar perspective (Todd & Akerstrom, 1987).

Another important finding that has been reported for planar surfaces, but which may not generalize to curved surfaces, involves the analysis of anisotropic blob textures. The left panel of Fig. 20 shows the image of a planar surface that is slanted in a vertical direction, with an anisotropic texture of randomly positioned ellipses that are all oriented in a diagonal direction. Although it is mathematically possible to determine the correct direction of slant in this image from an analysis of the texture gradients (Malik & Rosenholtz, 1997), any algorithm that is based on an assumption of texture isotropy would produce incorrect slant estimates in a diagonal direction. The empirical evidence indicates that human observers produce similar errors (Knill, 1998b; Rosenholtz & Malik, 1997), thus suggesting that the perceptual analysis of 3D shape from texture must be based at least in part on an implicit assumption that texture patterns are statistically isotropic. More recent research

<sup>6</sup> It is interesting to note that most naïve observers report that the apparent signs of curvature in these displays remain perceptually quite stable, despite the fact that the overall sign of relief is mathematically ambiguous. The most likely explanation of this phenomenon is that observers have strong biases to interpret ambiguous surfaces as viewed from above rather than from below (Mamassian & Landy, 2001; Reichel & Todd, 1990) or as globally convex rather than concave (Langer & Bühlhoff, 2001; Liu & Todd, 2004).

suggests, however, that this conclusion may not generalize to more complex surface geometries (Todd, Oomes, Koenderink, & Kappers, 2004). The right panel of Fig. 20 shows a complex doubly curved surface with a texture that is both anisotropic and inhomogeneous. Note in particular how the central region of this surface that appears to have a fronto-parallel orientation contains texture elements that are all noticeably elongated in a horizontal direction. This would not be possible if the perceptual interpretation were based on an assumption of texture isotropy. Indeed, Todd et al. (2004) found that observers' shape judgments for doubly curved surfaces with this type of texture are not detectably different from those that are obtained when the same surfaces are presented with isotropic textures.

There is other anecdotal evidence to suggest that the complexity of surface geometry can also influence the relative salience of different texture cues for the visual perception of 3D shape. Consider, for example, the pair of images of a hyperbolic cylinder that are presented in Fig. 21. The surface depicted in the left panel is textured with a pattern of contours that are all oriented in the direction of maximum curvature, which isolates the gradient of contour convergence (i.e., linear perspective). The image on the right, in contrast, shows the same surface with a pattern of contours that all oriented in the direction of minimum curvature, which isolates the gradient of contour compression. Note how the convergence gradient produces a compelling perception of surface slant, whereas the compression gradient does not. This phenomenon was first reported for planar surfaces over 35 years ago (Attneave & Olson, 1966; Gillam, 1968, 1970), and has more recently been extended to singly curved surfaces like the ones depicted in Fig. 21 (Li & Zaidi, 2000, 2001). One important caveat about convergence gradients, however, is that they generally require a relatively large amount of perspective in order to be perceptually effective. This is easily achieved with planar surfaces that can extend indefinitely in depth, but that is not the case for globally convex objects because of the effects of occlusion. To demonstrate this more clearly, the left panel of Fig. 22 shows a perspective projection of randomly deformed sphere that is depicted with a 20° field of view, and is textured with a series of planar cut contours (see Tse, 2002) that are all oriented in depth so as to isolate the gradient of contour convergence. Although this provides some minimal information about 3D shape, the overall magnitude of apparent depth is quite small. The right panel of Fig. 22 shows exactly the same object with exactly the same field of view, but in this case it is textured with a series of planar cuts that all have a fronto-parallel orientation so as to isolate the gradient of contour compression. Note that the perception of 3D shape is much more compelling in the right panel than in the left, which is the opposite of what occurs in Fig. 21 with singly curved surfaces.






Fig. 20. Perspective projections of a planar surface and a doubly curved surface with anisotropic textures.




Fig. 21. Perspective projections of a hyperbolic cylinder with contour textures in different orientations.




Fig. 22. Perspective projections of a randomly deformed sphere with contour textures in different orientations.

It is interesting to point out that the ability of human observers to correctly interpret the 3D shapes of the surfaces shown in the right panels of Figs. 20 and 22 is really quite remarkable, because these images violate the assumptions of every current model in the field. The only existing algorithm that can correctly interpret the optical projections of anisotropic blob textures is the one developed by Malik and Rosenholtz (1997). This algorithm works well with images of planar or singly curved surfaces, but it is not easily generalized for doubly curved surfaces. Thus, it would be difficult to

explain on the basis of that model why doubly curved surfaces with inhomogeneous anisotropic textures can be correctly interpreted by human observers, but that images of planar surfaces with homogeneous anisotropic textures appear systematically distorted. Similarly, existing computational analyses of contour textures typically assume that an observed surface is singly curved, and that its contours are constrained to lie along surface geodesics (Knill, 2001) or directions of principal curvature (Stevens, 1981), but none of those constraints are satisfied by the images in Fig. 22 (see also Todd & Reichel, 1990; Todd et al., 2004). It is not at all clear from the available empirical evidence how human observers are able to correctly interpret such theoretically anomalous texture patterns. These demonstrations suggest that observers may employ some heretofore unknown strategies for determining shape from texture on complex doubly curved surfaces, and the investigation of those strategies will remain as an intriguing problem for future research.

#### Acknowledgments

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

## References

- Aloimonos, J. (1988). Shape from texture. *Biological Cybernetics*, 58, 345–360.
- Attneave, F., & Olson, R. K. (1966). Inferences about visual mechanisms from monocular depth cues. *Psychonomic Science*, 4, 246–256.
- Blake, A., Bülthoff, H. H., & Sheinberg, A. (1993). Shape from texture: Ideal observers and human psychophysics. *Vision Research*, 33, 1723–1737.
- Blake, A., & Marinos, C. (1990). Shape from texture: Estimation, isotropy and moments. *Artificial Intelligence*, 45, 323–380.
- Bradshaw, M. F., Glennerster, A., & Rogers, B. J. (1996). The effect of display size on disparity scaling from differential perspective and vergence cues. *Vision Research*, 36, 1255–1264.
- Braunstein, M. L., & Payne, J. W. (1969). Perspective and form ratio as determinants of relative slant judgments. *Journal of Experimental Psychology*, 81, 584–590.
- Brown, L. G., & Shvaytser, H. (1990). Surface orientation from projective foreshortening of isotropic texture autocorrelation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 584–588.
- Buckley, D., Frisby, J., & Blake, A. (1996). Does the human visual system implement an ideal observer theory of slant from texture? *Vision Research*, 36, 1163–1176.
- Clerc, M., & Mallat, S. (2002). The texture gradient equation for recovering shape from texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24, 536–549.
- Cornilleau-Péres, V., Wexler, M., Droulez, J., Marin, E., Miège, C., & Bourdoncle, B. (2002). Visual perception of planar orientation: Dominance of static depth cues over motion cues. *Vision Research*, 42, 1403–1412.
- Cumming, B. G., Johnston, E. B., & Parker, A. J. (1993). Effects of different texture cues on curved surfaces viewed stereoscopically. *Vision Research*, 33(5–6), 827–838.
- Cutting, J. E., & Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, 113(2), 198–216.
- Davis, L. S., Janos, L., & Dunn, S. M. (1983). Efficient recovery of shape from texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5, 485–492.
- Frisby, J. P., & Buckley, D. (1992). Experiments on stereo and texture cue combination in human vision using quasi-natural viewing. In G. A. Orban & H. H. Nagel (Eds.), *Artificial and biological visual systems*. Berlin: Springer-Verlag.
- Gårding, J. (1992). Shape from texture for smooth curved surfaces in perspective projection. *Journal of Mathematical Imaging and Vision*, 2, 327–350.
- Gårding, J. (1993). Shape from texture and contour by weak isotropy. *Artificial Intelligence*, 64, 243–297.
- Gibson, J. J. (1950a). The perception of visual surfaces. *American Journal of Psychology*, 63, 367–384.
- Gibson, J. J. (1950b). *The perception of the visual world*. Boston: Houghton Mifflin.
- Gillam, B. (1968). Perception of slant when perspective and stereopsis conflict: Experiments with aniseikonic lenses. *Journal of Experimental Psychology*, 78, 299–305.
- Gillam, B. (1970). Judgments of slant on the basis of foreshortening. *Scandinavian Journal of Psychology*, 11, 31–34.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries and emergent segmentations. *Perception & Psychophysics*, 38, 141–171.
- Grossberg, S., & Mingolla, E. (1987). Neural dynamics of surface perception: Boundary webs, illuminants, and shape from shading. *Computer Vision, Graphics, and Image Processing*, 37, 116–165.
- Kanatani, K., & Chou, T. C. (1989). Shape from texture: General principal. *Journal of Artificial Intelligence*, 38, 1–48.
- Knill, D. C. (1998a). Discrimination of planar surface slant from texture: Human and ideal observers compared. *Vision Research*, 38, 1683–1711.
- Knill, D. C. (1998b). Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture. *Vision Research*, 38, 2635–2656.
- Knill, D. C. (2001). Contour into texture: Information content of surface contours and texture flow. *Journal of the Optical Society of America A*, 18, 12–35.
- Langer, M. S., & Bülthoff, H. H. (2001). A prior for local convexity in local shape from shading. *Perception*, 30, 403–410.
- Li, A., & Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research*, 40, 217–242.
- Li, A., & Zaidi, Q. (2001). Information limitations in perception of shape from texture. *Vision Research*, 41, 1519–1534.
- Li, A., & Zaidi, Q. (2003). Observer strategies in perception of 3-D shape from isotropic surfaces: Developable surfaces. *Vision Research*, 43, 2741–2758.
- Liu, B., & Todd, J. T. (2004). Perceptual biases in the interpretation of 3D shape from shading. *Vision Research*, 44, 2135–2145.
- Malik, J., & Rosenholtz, R. (1997). Computing local surface orientation and shape from texture for curved surfaces. *International Journal of Computer Vision*, 23, 149–168.
- Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, 41, 2653–2668.
- Phillips, R. J. (1970). Stationary visual texture and the estimation of slant angle. *Quarterly Journal of Experimental Psychology*, 22, 389–407.
- Purdy, W. C. (1958). *The hypothesis of psychophysical correspondence in space perception*. Unpublished Ph.D., Cornell University, Ithaca, NY.
- Reichel, F. D., & Todd, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 653–664.
- Rogers, B. J., & Bradshaw, M. F. (1995). Disparity scaling and the perception of frontoparallel surfaces. *Perception*, 24, 155–179.
- Rosenholtz, R., & Malik, J. (1997). Surface orientation from texture: Isotropy or homogeneity (or both)? *Vision Research*, 37, 2283–2293.
- Stevens, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence*, 17, 47–53.
- Super, B., & Bovik, A. (1995). Shape from texture using local spectral moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, 333–343.
- Tibau, S., Willems, B., Van Den Berg, E., & Wagemans, J. (2001). The role of the center of projection in the estimation of slant from texture of planar surfaces. *Perception*, 30, 185–193.
- Todd, J. T., & Akerstrom, R. A. (1987). The perception of three dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 242–255.
- Todd, J. T., & Oomes, A. H. (2002). Generic and non-generic conditions for the perception of surface shape from texture. *Vision Research*, 42, 837–850.
- Todd, J. T., Oomes, A. H. J., Koenderink, J. J., & Kappers, A. M. L. (2004). Perception of doubly curved surfaces from anisotropic textures. *Psychological Science*, 15, 40–46.
- Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, 96, 643–657.
- Todd, J. T., & Reichel, F. D. (1990). The visual perception of smoothly curved surfaces from double projected contour patterns. *Journal of*

- Experimental Psychology: Human Perception and Performance*, 16, 665–674.
- Tse, P. U. (2002). A contour propagation approach to surface filling-in and volume formation. *Psychological Review*, 109, 91–115.
- Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17, 17–45.
- Zaidi, Q., & Li, A. (2002). Limitations on shape information provided by texture cues. *Vision Research*, 42, 815–835.