PERGAMON

Vision Research 42 (2002) 837–850

Vision Research

www.elsevier.com/locate/visres

Generic and non-generic conditions for the perception of surface shape from texture

James T. Todd *, Augustinus H.J. Oomes

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

Received 29 January 2001; received in revised form 18 June 2001

Abstract

Li and Zaidi (Vision Research 40 (2000) 217; 41 (2001) 1519) have recently argued that there are two necessary conditions for the perception of 3D shape from texture: (1) the texture pattern must have a disproportionate amount of energy along directions of principal curvature; and (2) the surface must be viewed with a noticeable amount of perspective. In the present article we present evidence that these conclusions are only valid under a limited set of non-generic viewing conditions. Other relevant factors that need to be considered in this context include the distribution of curvature on an object's surface and the set of possible viewing directions from which it can be observed. For generic viewing directions and patterns of curvature, the perception of surface curvature from texture is only minimally affected by the orientation spectrum of the texture pattern or the amount perspective in its optical projection. Li and Zaidi (Vision Research 41 (2001) 1519) have also identified two characteristic patterns of image contours, which they claim to be the only possible source of information within textured images for determining the direction of surface slant or the sign of surface curvature. In the present article we attempt to show that these characteristic patterns can only arise in natural vision for a limited set of non-generic viewing directions. We also review several other factors that can influence the perceived direction of slant or the perceived sign of curvature, which have been identified previously by other investigators. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Perception; Shape; Texture

1. Generic and non-generic conditions for the perception of surface shape from texture

The concept of optical texture is a relatively new one within perceptual psychology. It was first introduced as a possible source of information about 3D surface shape by Gibson (1950a) in his influential book *The perception of the visual world*. Gibson's arguments were supported primarily by the numerous examples he provided from drawings and photographs, but they also inspired a new program of empirical research and computational modeling that continues today. Most of these subsequent investigations have focused on the perceived slants of planar surfaces (e.g. Attneave & Olson, 1966; Gibson, 1950b; Knill, 1998a,b,c; Phillips, 1970; Rosenholtz & Malik, 1997), though there have been several recent studies that have also examined the perception of surface curvature from texture (Buckley & Frisby, 1993; Cumming, Johnston, & Parker, 1993; Cutting & Millard, 1984; Knill, 2001; Reichel & Todd, 1990; Todd & Akerstrom, 1987; Todd & Reichel, 1989, 1990).

One of the latest contributions to this topic has been provided by Li and Zaidi (2000, 2001), who investigated how the spectral properties of a texture can influence the perceived structure of sinusoidally corrugated surfaces. The right panel of Fig. 1 shows a textured image that is adapted from their study. The surface was created using 3D Studio Max by Kinetix to conform with the specifications described in their paper, and it was parametrically textured with a sinusoidal plaid pattern using the Darktree texture plug-in by Darktree Studios. Li and Zaidi investigated the perceived structure of this pattern using a local relative depth probe task similar to those used in previous studies (e.g. see Reichel & Todd, 1990; Reichel, Todd, & Yilmaz, 1995; Todd & Reichel, 1989). This technique is quite useful for measuring the perception of surface concavities and convexities, though it is insensitive to the magnitude of perceived surface relief

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

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

PII: S0042-6989(01)00234-6

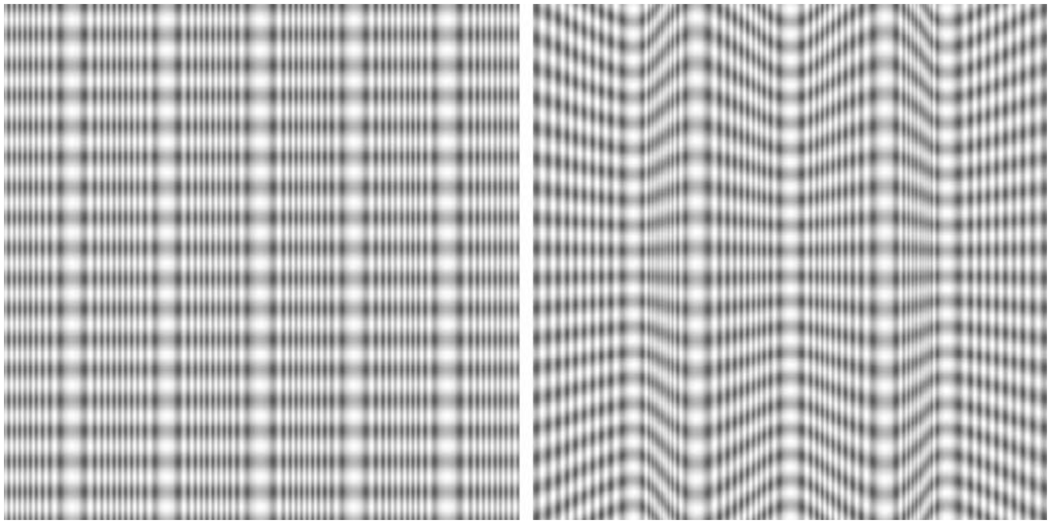

Fig. 1. A sinusoidally corrugated surface with a sinusoidal plaid texture, similar to the one used by Li and Zaidi (2000). The depicted surface has a period of 5.13 cm and a peak-to-trough amplitude of 19 cm. The image on the left is shown under orthographic projection. The one on the right is a perspective projection from a simulated viewing distance of 100 cm and a visual angle of 8.8°.

(see, however, the more recent refinements of Koenderink, van Doorn, & Kappers, 1996; Koenderink, van Doorn, Kappers, & Todd, 2001). Within that limitation, Li and Zaidi found that the qualitative structure of this pattern is perceived veridically.

In order to determine the relative contributions of different spectral components on observers' perceptions, Li and Zaidi produced several variations of this stimulus, including ones in which the horizontal and vertical components were presented in isolation (see Fig. 2). Their results revealed that surfaces with horizontal grating textures are perceived quite accurately, but that vertically oriented textures do not support a strong impression of 3D structure. They also created a variety of other textures whose orientation and frequency spectra were systematically manipulated. Based on observers' judgments of 3D structure from these textures, they reached the following conclusion: "... we find that three-dimensional shape is conveyed only by textures that contain variations in a direction orthogonal to the direction of the corrugation, but that this component must be detectable, which happens when either the global frequency spectrum of the texture pattern is discrete in orientations, or is elongated with an aspect ratio of 5 or greater in favor of the critical orientation."

It is important to recognize when considering this conclusion that the authors are not referring to the frequency spectrum for an image of a corrugated surface, like the ones shown in Figs. 1 and 2. They are referring instead to the spectrum of a planar texture before

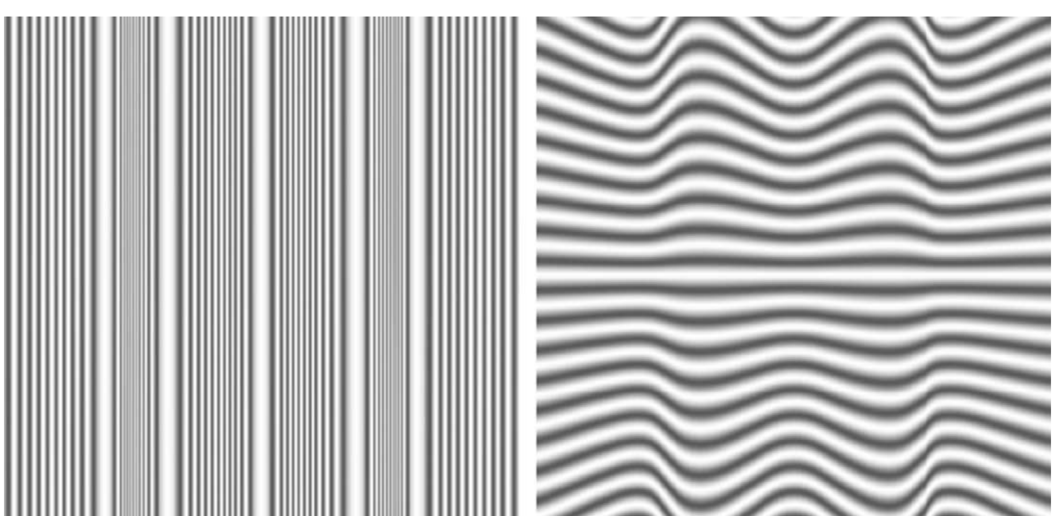

Fig. 2. The same surface as in the right panel of Fig. 1 with a sinusoidal grating texture in both a vertical and a horizontal orientation.

it is mapped onto a 3D surface. For example, the left panel of Fig. 3 shows the sinusoidal plaid texture that was used to create the images in Fig. 1, and the right panel of this figure shows the global frequency spectrum of that texture. Note in this case that the energy is confined to just four discrete points. According to Li and Zaidi, it is the two along the vertical axis (i.e. the ones that represent the horizontal component of the plaid) that provide perceptually relevant information about 3D structure for a surface that is corrugated in the horizontal direction.

A specific hypothesis about how this information is manifested in an image of a corrugated surface is provided in Li and Zaidi (2001). Note in the right panel of Fig. 2 that the image contours in convex surface regions curve inward toward the center of the display, whereas those in concave regions curve outward. According to Li and Zaidi, these inward and outward configurations provide the proximal stimulus information for observers' perceptions of concavity and convexity. They also performed a mathematical analysis to investigate how the pattern of contour curvatures in an image is influenced by the orientation of texture contours on a surface. Their results revealed that the inward and outward configurations that distinguish concavities from convexities only occur if the texture contours are oriented within a few degrees of the axis of maximum curvature, which could explain why those components of the texture are so critical for the visual perception of 3D shape.

Li and Zaidi (2001) acknowledge that non-critical texture orientations can sometimes interact with the critical ones to influence the perceptual appearance of 3D structure. Note, for example, that the right panels of Figs. 1 and 2 do not have the same apparent shape, even though their critical components are identical. When the non-critical components are presented in isolation, however, Li and Zaidi (2001) argue that "... observers utilize this information, at best, to distinguish curvatures from slants, which is all that is supported by this information."

Another important factor that can influence the perception of shape from texture is the amount of perspective in an image. Perspective is typically measured as a ratio from zero to one between an object's extension in depth and its farthest distance from the observer (see Braunstein, 1976). Images that have zero perspective are sometimes referred to as orthographic projections. Although this can never be achieved perfectly for curved surfaces in natural vision, it is closely approximated when an object is viewed from a relatively far distance. In an effort to demonstrate the importance of perspective for human perception, Li and Zaidi (2000) investigated observers' perceptions of a corrugated surface with a sinusoidal plaid texture that was presented under orthographic projection. A similar image that is adapted from their study is shown in the left panel of Fig. 1. Note that the depicted surface in that image appears completely flat with no apparent depth at all. Li and Zaidi concluded from this observation that perspective projections are a necessary condition for the visual perception of 3D shape from texture.

Before we go on to comment about the validity of these conclusions, it is important to point out that there is some ambiguity about the range of phenomena for which they are intended to apply. On some occasions, Li and Zaidi argue that perspective projections of contours oriented along lines of maximum curvature are necessary for a veridical perception of the sign of surface curvature or slant. More commonly, however, they express their conclusions in a more general manner, as in the quotation cited above, where it is stated that the critical orientations are necessary for the perception of

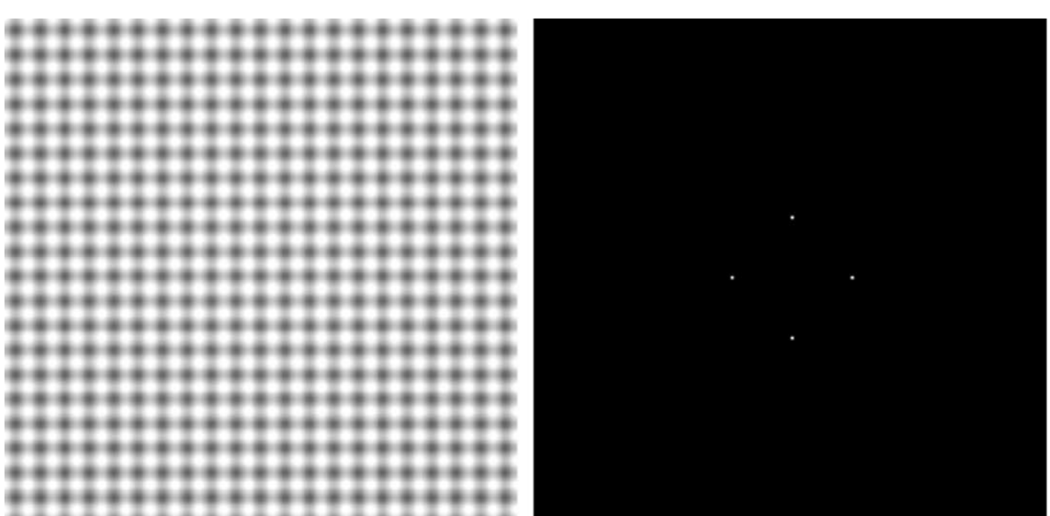

Fig. 3. The sinusoidal plaid texture that was used to create Fig. 1, and its global Fourier frequency spectrum.

3D shape. If such statements are taken literally, it would imply that the critical information they have identified is important for the appearance of 3D structure, not just the sign of curvature or slant. This literal interpretation is reinforced by the stimuli they have employed in their experiments. In general, the textures they have used that do not yield veridical ordinal depth judgments also do not produce a compelling perception of a 3D surface (e.g. see the left panels of Figs. 1 and 2). In the present discussion, we will consider both of these aspects of 3D shape from texture. First we will examine a variety of factors that can influence the appearance of 3D structure and the magnitude of perceived surface relief, and then we will consider how the apparent sign of relief is perceptually determined.

What prompted us to write a commentary about Li and Zaidi's experiments is that several of their findings appear to be inconsistent with other results that have been published in the literature. Consider, for example, the surface depicted in the right panel of Fig. 2 with horizontal contours in the direction of maximum curvature. The impression we get from this pattern is that horizontal cross-sections through the center of the surface appear nearly flat, and that the apparent amplitude of the corrugations increases gradually moving upward or downward from the center. Some of our observers report that this display is also perceptually multi-stable, such that the corrugated ridges at the top and bottom can sometimes appear to be in counter-phase. It is interesting to note when considering these phenomena that similar surfaces with similar contour textures have been investigated by other researchers (e.g. see Knill, 2001; Stevens, 1981a; Todd & Reichel, 1989, 1990), but the stimuli in those earlier studies do not produce the apparent variations in surface amplitude or the perceptual multi-stability that are evident in Fig. 2. This could be an indication that the stimulus parameters used by Li and Zaidi may constitute a special case, which somehow excludes potential sources of information about 3D shape from texture that are available in other contexts.

Another important aspect of Li and Zaidi's experiments that we believe should be given a more detailed examination involves the importance of perspective for the perception of 3D shape from texture. Although perspective is clearly necessary to perceive 3D structure for the particular displays employed in their experiment, there have been numerous published demonstrations in the literature to show that orthographic projections of textured surfaces can, under appropriate conditions, provide perceptually compelling information about surface curvature (e.g. see Knill, 2001; Stevens, 1981a; Todd & Reichel, 1989, 1990). As we will attempt to show in the discussion that follows, the conditions for which perspective is necessary to perceive 3D shape from texture are highly non-generic.

During our efforts to reproduce the stimuli used by Li and Zaidi, we have informally investigated a wide variety of textures for the corrugated surface shown in Fig. 1, and the results have further fueled our suspicions about the generality of their conclusions. One of the most common types of texture that has been used by researchers in this area since the 1950s involves random patterns of polka dots. A typical example of this type of texture is shown in the left panel of Fig. 4, and its global Fourier frequency spectrum is shown in the right panel. Note that this texture has roughly uniform energy in all possible directions. Thus, according to Li and Zaidi, it does not provide the necessary information that is needed to perceive 3D shape from texture. When we apply this texture to the same corrugated surface depicted in Fig. 1, the results appear to confirm their prediction (see Fig. 5), because the appearance of 3D structure is quite

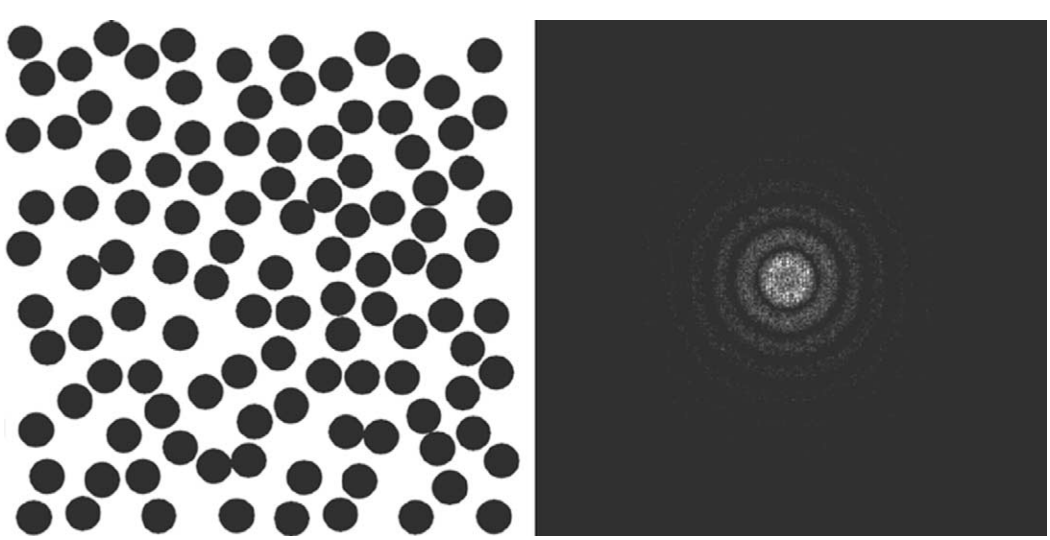

Fig. 4. The polka dot texture that was used to create Fig. 5, and its global Fourier frequency spectrum.

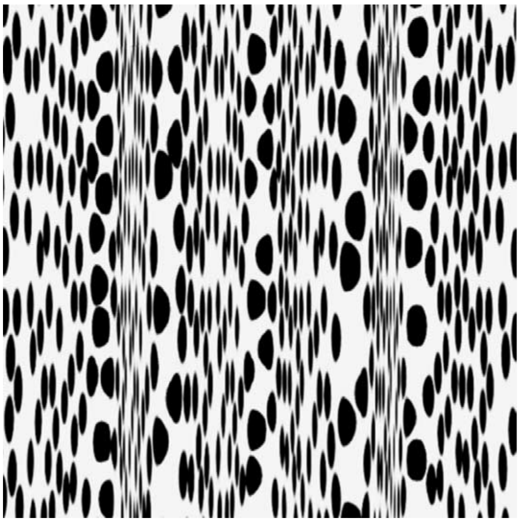

Fig. 5. The same surface as in the right panel of Fig. 1 with a random polka dot texture.

weak. The problem with this observation, however, is that it is not consistent with several other studies that have been reported previously in the literature, in which textures similar to the one in Fig. 4 have been shown to provide perceptually salient information about 3D surface structure (e.g. Buckley & Frisby, 1993; Cumming et al., 1993; Cutting & Millard, 1984; Phillips, 1970; Todd & Akerstrom, 1987). These findings provide additional evidence that there is something unusual about the particular surface configuration used by Li and Zaidi, and that their results may not generalize to other contexts.

To better understand the unusual nature of their stimuli, it is useful to consider two general factors that can produce gradients of texture within a visual image. One of these factors is the amount of perspective, which produces systematic changes in the projected sizes of optical texture elements. For images of an extended ground plane, the perspective ratio can be close to one, and the projected sizes of near elements will be many times larger than those farther away in depth. For most other solid objects, however, the amount of perspective that can be achieved in practice is quite small, except under extraordinary viewing conditions. For the stimuli used by Li and Zaidi the perspective ratio was 0.19, so that the relative projected sizes of near and far elements differed by about 19%. For surfaces presented under orthographic projection these size variations are eliminated altogether.

The other general factor that produces gradients of texture involves changes in surface orientation relative to the line of sight. As texture elements become more and more slanted relative to the viewing direction, their projected shapes become systematically compressed—a phenomenon that is often referred to in the literature as foreshortening. Note that changes in surface orientation due to curvature can produce gradients of foreshortening even when images are viewed under orthographic projection. Moreover, for surfaces with well-defined occlusion contours, the range of foreshortening will vary over the full range between zero and one regardless of the type of projection employed. Thus, for images of curved surfaces, the available information from foreshortening gradients will generally be greater than the information provided by gradients of texture size (see Cutting & Millard, 1984; Stevens, 1981b; Todd & Akerstrom, 1987).

Let us now examine the specific viewing geometry used by Li and Zaidi, which is shown schematically in Fig. 6. The stimuli in their study depicted sinusoidally corrugated surfaces with a period of 5.13 cm and a peak-to-trough amplitude of 19 cm, presented at a viewing distance of 100 cm. An important thing to note when examining this figure is that the depicted surface has negligible curvature over most of its extent, except within small neighborhoods at the peaks and troughs where the surface orientation changes abruptly. In addition, the local optical slant over much of the surface is close to its maximum value of 90°. It is because of this unusually severe depth profile that that the circle textures of Fig. 5 cannot adequately specify the depicted surface structure. The foreshortening gradients at the peaks and troughs are much too high relative to the scale of the texture to be perceptually useful, and the foreshortening gradients in

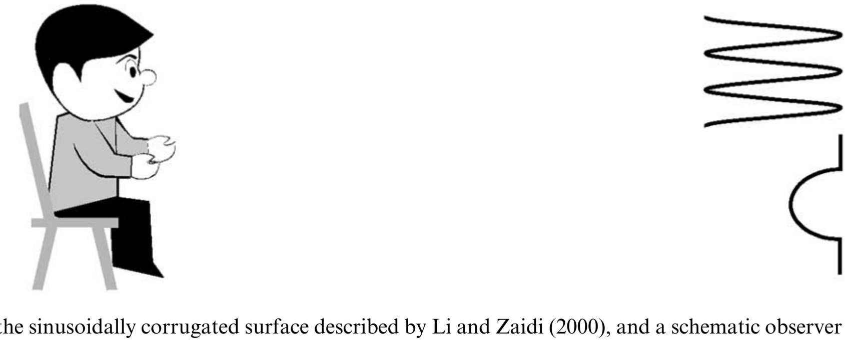

Fig. 6. A cross-section of the sinusoidally corrugated surface described by Li and Zaidi (2000), and a schematic observer at the same viewing distance used in their study. Just below the sinusoidal surface is the cross-section of an elliptical cylinder that is depicted in Figs. 7–12.

the planar regions are too small. This is the first of several degenerate properties in their displays that limited the available information.

2. Degenerate surface structures

Surfaces that do not have gradual orientation changes relative to the viewing direction are degenerate for providing information about 3D shape from gradients of texture compression. To demonstrate the importance of gradual orientation changes it is useful to compare the sinusoidally corrugated surface used by Li and Zaidi with an elliptical cylinder whose cross-section is shown just beneath it in Fig. 6. The width of this surface (including its background) is 15.4 cm and its amplitude in depth is 7 cm. The right panel of Fig. 7 shows an image of this surface with a polka dot texture that was generated from a simulated viewing distance of 20 cm, i.e. the perspective ratio for this image is 0.35 and the simulated viewing angle is 40°. The left panel of this figure shows the same surface under orthographic projection. Because this surface has more gradual orientation changes than the one depicted in Fig. 5, the projected distortions of the polka dots provide perceptually compelling information about the overall pattern of surface curvature.

It is important to recognize that the polka dot textures employed in these figures do not contain discrete-oriented energy in the directions of principal curvature nor do they have a frequency spectrum that is elongated in the critical directions. Nevertheless, the magnitude of perceived depth in these figures appears much larger than for the surface depicted in Fig. 1, which does contain discrete oriented energy in the directions of principal curvature and has a simulated depth that is more than twice as large (see Fig. 6). It is also important to note in this regard that the apparent 3D structures for both images in Fig. 7 are quite similar, thus indicating that perspective projections are not a necessary condition for the perception of 3D shape from texture. Although perspective can influence the magnitude of perceived depth in these displays, its effect is quite small relative to other stimulus factors (Todd & Akerstrom, 1987).

Let us now examine how the perception of this surface is affected by contour textures whose global frequency spectra contain energy in just one orientation. Fig. 8 shows our elliptical cylinder with a random contour texture that is oriented in the horizontal direction. The right panel of this figure shows the surface with an exaggerated perspective, as it would appear from a viewing distance of 20 cm with a visual angle of 40°. The left panel shows the surface from a viewing distance of 100 cm with a visual angle of 8.8° – the same viewing geometry used by Li and Zaidi. Note in this case that the magnitude of perceived relief is dramatically influenced by the simulated viewing distance. Indeed, when this same object is presented under orthographic projection, the appearance of 3D form is eliminated altogether, and the image is perceived as a 2D pattern of parallel bars. Why does perspective have a greater effect on the perception of shape from contour textures in Fig. 8 than it does on the perception of polka dot textures in Fig. 7? As it turns out, this strong effect of perspective is dependent on a combination of two additional degenerate conditions, whose mutual co-occurrence in natural vision is highly non-generic.

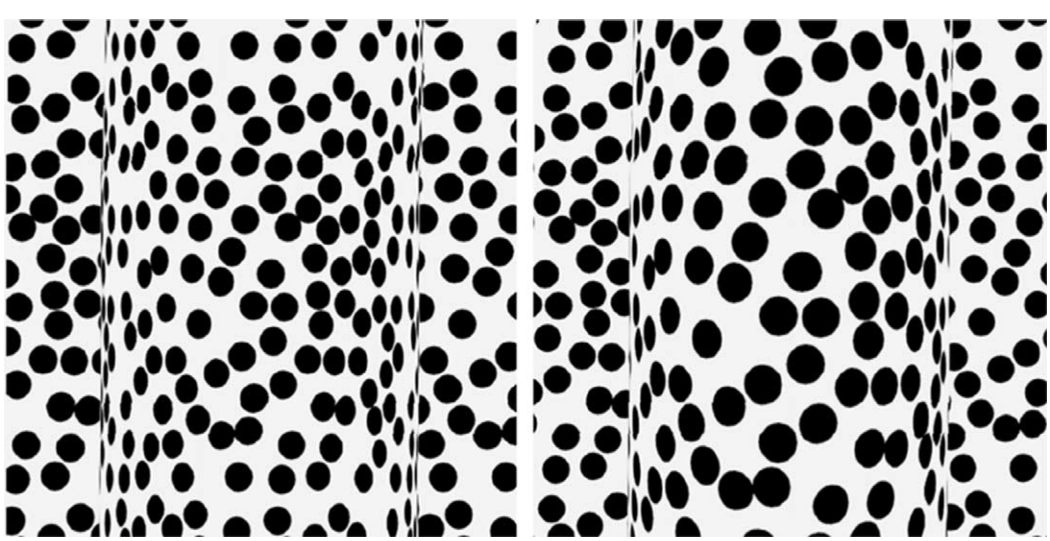

Fig. 7. An elliptical cylinder with a polka dot texture. The image on the left is shown under orthographic projection. The one on the right is a perspective projection from a simulated viewing distance of 20 cm and a visual angle of 40°. Note that the variation in perspective between these images has only a minimal effect on the perceived 3D structure.

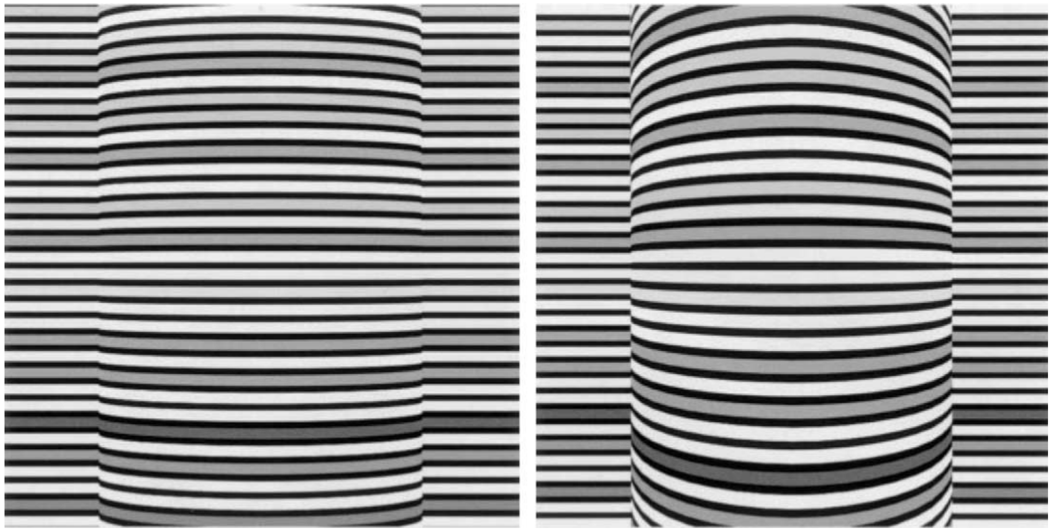

Fig. 8. An elliptical cylinder with a horizontally oriented contour texture. The image on the left is a perspective projection from a simulated viewing distance of 100 cm and a visual angle of 8.8°—the same conditions used by Li and Zaidi (2000). The one on the right is a perspective projection from a simulated viewing distance of 20 cm and a visual angle of 40°. Note that the effects of perspective in this case are quite large.

3. Degenerate texture orientations

The left panel of Fig. 9 shows a contour texture similar to the one that was used to generate Fig. 8, but in this case the contours have been rotated 45° to a diagonal orientation. The global frequency spectrum for this texture is shown in the right panel. Note that all of the energy is oriented in a diagonal direction, and that there is no energy at all along the directions of principal curvature. In order to assess the effect of this texture on the visual perception of 3D shape, we have applied it to an elliptical cylinder in Fig. 10, using both orthographic and perspective projections. Note that both of these images provide a clear perceptual impression of a curved surface in depth. It is clear from this demonstration that diagonal contour textures can provide perceptually salient information about the structure of curved surfaces (see also Knill, 2001). Under certain viewing conditions, moreover, diagonal contour textures can be much more informative than those oriented in the direction of maximum curvature—the direction that Li and Zaidi claim is essential for the visual perception of 3D shape from texture.

The strong perceptual differences between the images in Figs. 8 and 10 have a simple theoretical explanation. What makes these images appear three-dimensional is that the curvature of the image contours is perceptually attributed to the structure of the depicted surface. In any

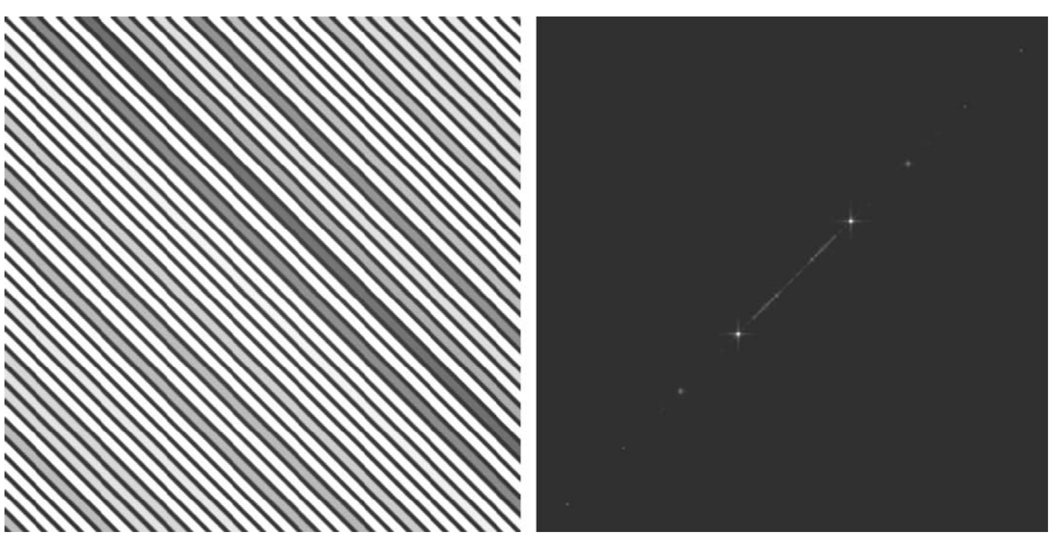

Fig. 9. The diagonal contour texture that was used to create Fig. 10, and its global Fourier frequency spectrum.

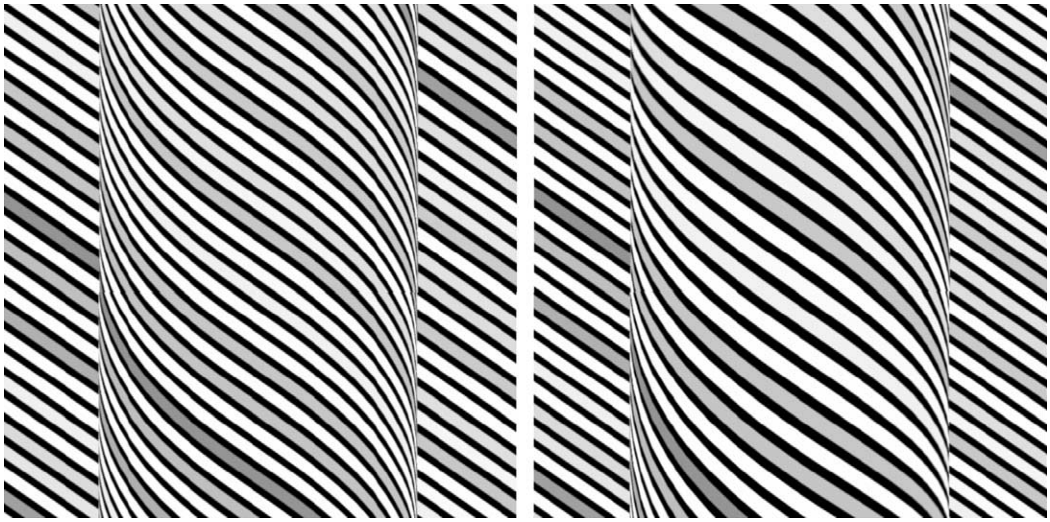

Fig. 10. An elliptical cylinder with a diagonally oriented contour texture. The image on the left is shown under orthographic projection. The one on the right is a perspective projection from a simulated viewing distance of 20 cm and a visual angle of 40°. When the contours are oriented in a diagonal direction, it eliminates the strong effect of perspective that occurs with horizontal contours in Fig. 8.

situation where the image contours have little or no curvature, then the depicted surface is perceived as flat. Note in Figs. 8 and 10 that the apparent 3D curvature of the surface varies systematically with the 2D curvatures of its image contours in each local region. The left panel of Fig. 8 appears relatively flat because its contours have relatively little curvature. How might this occur in natural vision? The most likely possibility, of course, is that the observed surface really is flat. Alternatively, it could be a curved surface in a non-generic pose such that its contours are only curved in a direction that is parallel to the observer's line of sight. It is important to note that there are three critical conditions that must be satisfied simultaneously in order to produce this degenerate special case: First, the surface contours in three-dimensional space must be planar space curves; second, the viewing direction must be parallel to those planes; and third, the magnitude of perspective must be small.

Using the texture mapping procedures employed thus far, the first requirement to create this special case can only occur if the contours are oriented in a direction of principal curvature. When 2D wall paper patterns are mapped homogeneously onto a singly curved surface, any straight contour in the texture will always lie along a surface geodesic. In directions of principal curvature, the surface geodesics are planar space curves. If viewed end on with small amounts of perspective, their optical projections will form a pattern of straight lines (e.g. see the left panel of Fig. 1). However, if the contours are oriented in any direction other than the axes of principal curvature (e.g. see Fig. 10), then they will form a pattern of helical space curves, whose optical projections will be curved regardless of the viewing direction or the magnitude of perspective. That is why the left panel of Fig. 10 appears much more curved than does the left panel of Fig. 8.

4. Degenerate viewing directions

The second requirement to create this special case is that the object must be viewed from a particular orientation. Suppose, for example, that a singly curved object with gradual orientation changes is textured with contours in the direction of maximum curvature, so that they form a pattern of planar cuts on its surface. Even within that restrictive context, the optical projections of those contours will still be generically curved if the object is viewed from any direction that is not parallel to the contour planes. In order to illustrate the importance of viewing direction in this situation, it is useful to compare the two images presented in Fig. 11. The left panel of this figure is based on a demonstration first published by Stevens (1981a). It depicts a sinusoidally corrugated surface under orthographic projection, whose amplitude is 85% smaller than for the surfaces shown in Figs. 1 and 2. The right panel shows the same surface under perspective projection, and from a different orientation that is parallel to the contour planes. The simulated viewing distance used to generate that image was 20 cm, and the simulated viewing angle was 40°. Despite its exaggerated perspective, however, the image in the right panel appears much less compelling than the one on the left that was produced under orthographic projection from a non-degenerate viewing direction.

In all of the images we have presented thus far, the surfaces were textured using a mathematical procedure that is analogous to covering a surface with wall paper.

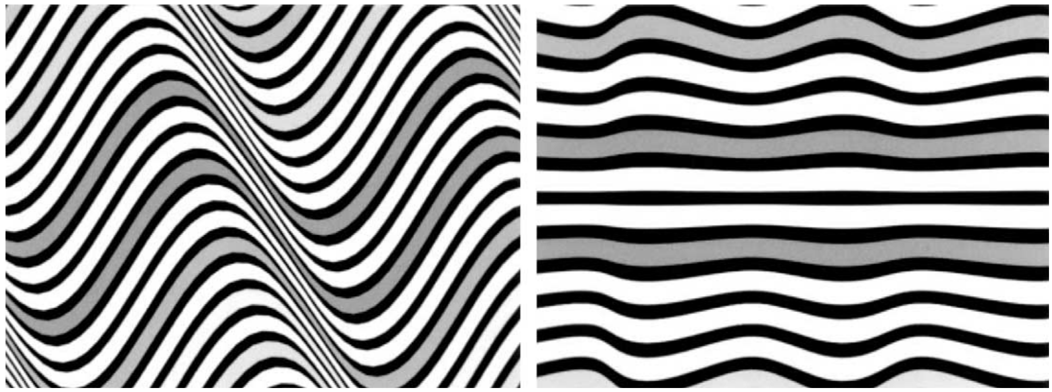

Fig. 11. A sinusoidally corrugated surface with a contour texture that is oriented in the direction of maximum curvature, so that its contours form a pattern of planar space curves. This surface has a period of 5.13 cm and a peak-to-trough amplitude of 2.85 cm. The image on the left is shown under orthographic projection at an oblique angle to the contour planes. The one on the right is a perspective projection from a degenerate viewing direction that is parallel to the contour planes at a distance of 20 cm and a visual angle of 40°.

An important limitation of this technique is that it is only appropriate for singly curved surfaces, such as cylinders, that are intrinsically planar. For doubly curved surfaces, it is not mathematically possible to map a wall paper pattern without tearing or distorting it. An alternative method of texturing such surfaces is to employ a volumetric texture map, which is analogous to sculpting an object out of a solid material. If that material is composed of compressed layers, such as wood or marble, then the pattern of surface reflectance visible on its surface will contain a systematic pattern of roughly parallel contours. However, whereas the contours on wallpaper textures lie along surface geodesics (see Figs. 8, 10, and 11), those that are created with volumetric textures of layered materials have contours that are confined to planar cuts through a surface (see Tse, 2001). An interesting special case for which the two procedures are equivalent is the one shown in Figs. 8 and 11, where contours on a singly curved surface are oriented in the direction of maximum surface curvature. That is the only possible situation where contours can lie along surface geodesics and also be confined to fixed plane.

The projected images of volumetric contour textures can provide compelling information about 3D surface shape, but only when a surface is viewed from a non-degenerate orientation that is sufficiently slanted relative to the planar cuts. The left panel of Fig. 12 shows a typical example that is adapted from the earlier demonstrations of Todd and Reichel (1989, 1990). It depicts a radial cosine surface with a peak-to-trough amplitude of 3 cm viewed under orthographic projection. Note that the contours form planar cuts through the surface and that they are not aligned with the directions of principal curvature or surface geodesics. The right panel of Fig. 12 shows the same surface from a degenerate orientation at a simulated viewing distance of 20 cm. Although this pattern has a strong polar perspective, the 3D structure of the depicted surface is barely discernable.

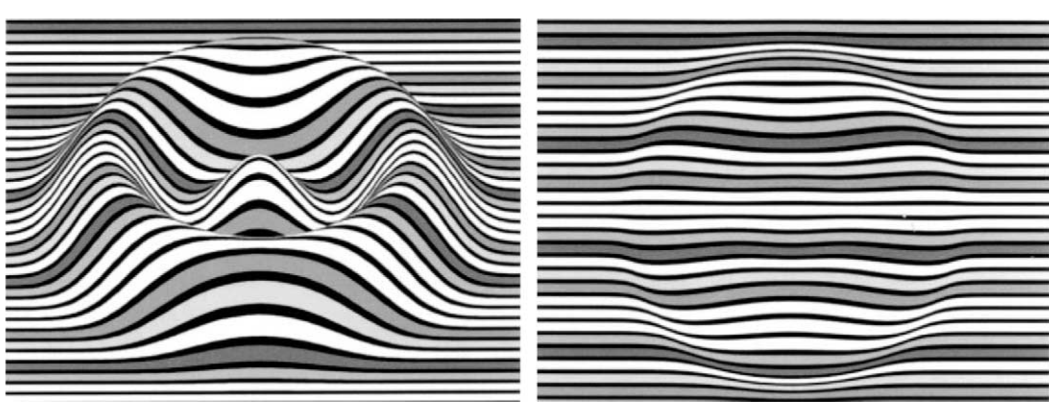

Fig. 12. A radial cosine surface with a 3D volumetric texture composed of parallel planar layers similar to that of marble or wood. The image on the left is shown under orthographic projection at an oblique angle to the planar layers. The one on the right is a perspective projection from a degenerate viewing direction that is parallel to the planar texture layers at a distance of 20 cm and a visual angle of 45°.

5. Concavities and convexities

A particularly interesting aspect of textured surfaces under orthographic projection is that their signs of curvature may be mathematically ambiguous, such that the image of a convex surface can be identical to one that is concave. Given that ambiguity, how is it possible that the surfaces presented under orthographic projection in Figs. 7, 10–12 can have perceptually stable curvatures? The most likely explanation of this stability is that observers' perceptions are biased to prefer one possible interpretation over another. For the objects depicted in Figs. 7 and 10, for example, there appears to be a bias to perceive the surfaces as convex, even though a concave interpretation would be equally valid. Similarly, in Figs. 11 and 12, there is a bias to perceive an overall surface slant, such that regions depicted near the bottom of the image appear closer in depth than those near the top (see Mamassian & Landy, 1998; Reichel & Todd, 1990). If those images are viewed upside down, the apparent sign of relief can be reversed. Both of these biases are reasonable expectations in the ecology of natural vision. Because of the topological and physical constraints on objects in the natural environment, convex surface regions are much more common than concave surface regions, and ground surfaces are more common than ceiling surfaces. We suspect that these biases are ever-present in the visual perception of shape from texture, though they can be overridden under certain conditions when the appropriate information is available.

Li and Zaidi (2001) have argued that "texture components within a few degrees of the axis of maximum surface curvature are the only components that form patterns that are distinct for different signs of curvatures and slants, and that observers utilize this information to correctly identify these surface shapes." The distinct image patterns to which they are referring are shown in Fig. 13. The image on the right depicts a convex circular cylinder with a diameter of 20 cm, a viewing distance of 50 cm and a visual angle of 40°. The image on the left shows the same cylinder with its front half removed, so that the visible surface has a concave curvature. According to Li and Zaidi, image contours that bow in opposite directions toward each other provide the critical information to specify a convex surface region, whereas those that bow in opposite directions away from one another provide the critical information to specify a concave surface region.

Li and Zaidi (2001) have correctly pointed out that these characteristic bowing patterns can only occur with perspective projections and for a limited range of texture orientations. Another important restriction they did not consider, however, is that these patterns are also limited to a relatively small range of non-generic viewing directions, for which the axis of zero curvature is approximately perpendicular to the observer's line of sight. Fig. 14 shows the same two surfaces depicted in Fig. 13, but from a viewing direction that is slanted 45° relative to the central axis of the cylinder. Note in this case that all of the contours in both images are curved in the same direction. Given the severe restrictions under which inward and outward bowing patterns of image contours can arise in natural vision, Li and Zaidi's assertion that these patterns are the only possible source of information for distinguishing concavities from convexities seems difficult to justify. It is also difficult to justify in light of other possible sources of information that have been identified previously by other investigators.

Consider, for example, the two curved surfaces depicted in Fig. 14. Some obvious sources of information for determining the direction of slant in these images include the gradient of compression in the projected contour widths, and the gradient of convergence in the outer boundaries of the figures (e.g. see Attneave & Olson, 1966; Gillam, 1968, 1970). Concavities and convexities can be distinguished in this context by how the contours are bowed relative to the directions of these gradients. For the concave surface, they are bowed in the direction of increasing compression (or convergence), and for the convex surface they are bowed in the

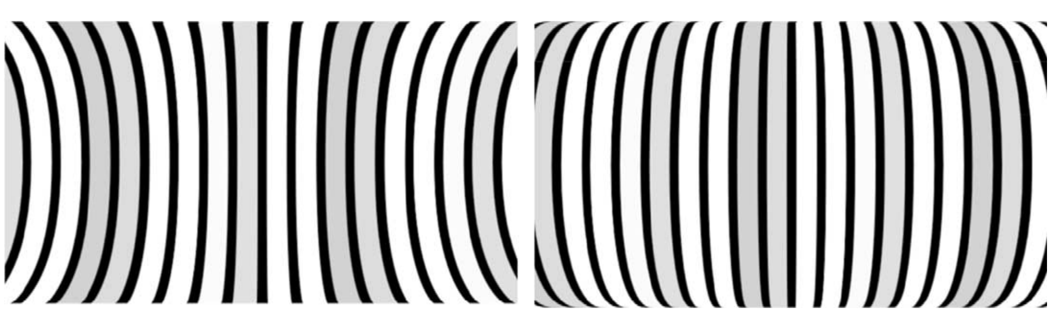

Fig. 13. Two images of a circular cylinder with a diameter of 20 cm, a viewing distance of 50 cm and a visual angle of 40°. The one on the right shows an outside view in which the surface is convex, whereas the one on the left shows an inside view in which the surface is concave. Note that the contours on the right and left of each figure are bowed in opposite directions. For the convex surface they bow toward each other, and for the concave surface they bow away from one another. According to Li and Zaidi (2001), this is the only aspect of textured images by which it is possible to distinguish concavities from convexities.

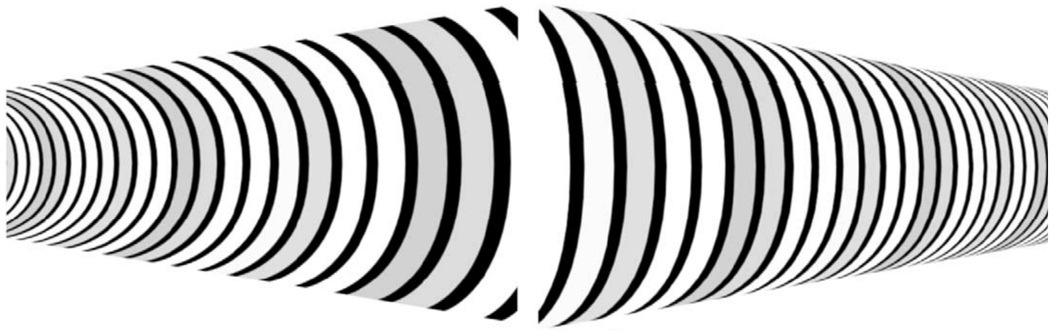

Fig. 14. The same two surfaces depicted in Fig. 13 from a viewing direction that is slanted 45° relative to the central axis of the cylinder. Note in this case that all of the contours are bowed in the same direction for both the convex surface on the right and the concave surface on the left.

opposite direction. It is especially interesting to note in this example how potential information about surface relief can interact with an observer's perceptual biases. For the image on the right, the texture information is consistent with the natural bias to prefer convexities over concavities, and that image has a stable perceptual appearance as a convex cylinder that is slanted in depth. For the image on the left, in contrast, the texture information conflicts with the bias, and the resulting appearance is perceptually multi-stable. It can appear as a concave circular cylinder slanted in depth or as a convex circular cone that is slanted in the opposite direction.

Gradients of texture size provide another potential source of information about the direction of surface slant or the sign of surface curvature that have been studied extensively since the 1950s (e.g. see Purdy, 1958; Stevens, 1981b). These gradients are most informative for surfaces such as ground planes that are visible over a large distance to produce the highest possible levels of perspective. There have been numerous demonstrations published in the literature to show how gradients of texture size provide veridical information about the direction of surface slant (e.g. see Gibson, 1950b; Rosenholtz & Malik, 1997), and these demonstrations can easily be extended for the perception of surface curvature. The right panel of Fig. 15 shows a variant of the polka dot planar surfaces that are commonly presented in introductory textbooks to exemplify how texture gradients provide information about surface slant. The surface in this case has a terraced structure with two concave regions and one that is convex. Although this texture does not contain salient contours in the critical directions identified by Li and Zaidi (2000, 2001), the gradients of texture size and shape provide sufficient information for observers to perceive the correct signs of curvature. These gradients can also be exploited to

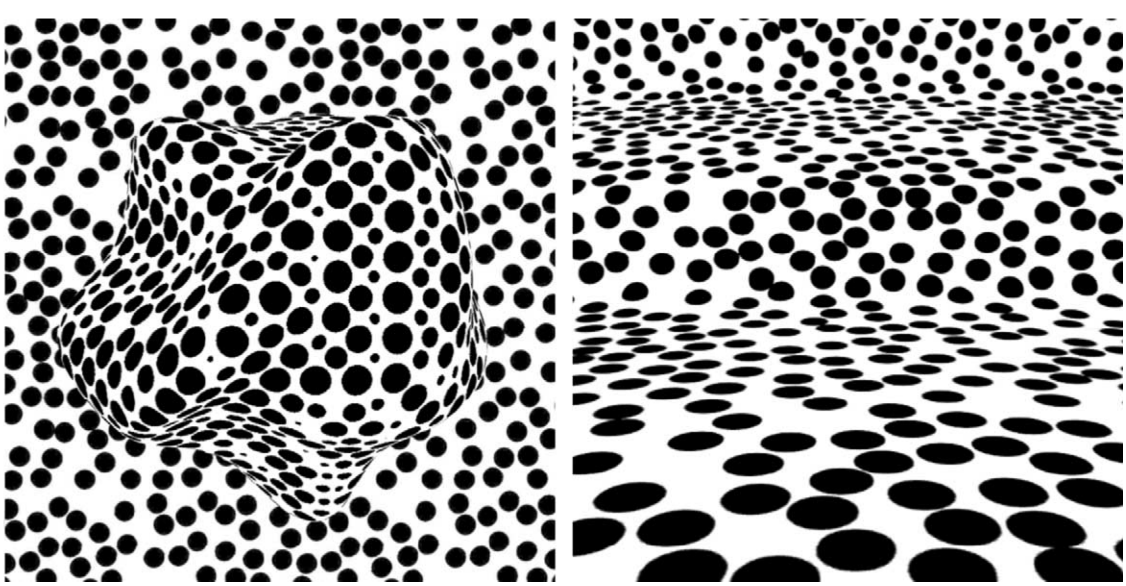

Fig. 15. Two curved surfaces with polka dot textures under strong perspective projection. The one on the left was carved out of a volume of black spheres. The one on the right has a 2D wall paper texture. Although neither of these textures has perceptually salient contours along directions of principal curvature, it is still possible to correctly perceive the sign of surface curvature and slant.

accurately determine the patterns of concavities and convexities on doubly curved surfaces, as is demonstrated in the left panel of Fig. 15.

There is one source of information in textured images about the sign of surface relief that is available even under orthographic projection. For generic views of most natural objects some surface regions will be occluded by others. Koenderink and van Doorn (1982) and Koenderink (1984) have shown mathematically that the occlusion contours on smoothly curved surfaces provide potential information about the sign of curvature in their immediate local neighborhoods. The importance of occlusion contours for the visual perception of shape from texture has been investigated by Reichel and Todd (1990), who also studied how this information interacts with the perceptual bias to prefer ground surfaces over ceiling surfaces. The left panel of Fig. 12 is similar to the stimuli employed in their experiment. If this image is viewed upside down the occlusion information will be incompatible with a ground surface interpretation. What can happen in that case is that the perceived relief will be determined by the occlusion information in some regions and by a ground surface bias in others. When this occurs, it produces a piece-wise inversion in depth that is perceptually disturbing in much the same way as an impossible figure.

It is important to keep in mind when considering the effects of smooth occlusion contours that gradients of texture provide a primary source of information by which these occlusions are perceptually specified. Note in Figs. 12 and 15, for example, that the occlusion contours are revealed by abrupt gradients of texture orientation or shape, and that the attached sides of these occlusions are indicated by a high degree of texture compression or foreshortening. For images such as these, the appearance of occlusions is a fundamental aspect of the perception of shape from texture that is of comparable importance to the appearance of surface slant or curvature.

6. Summary and conclusions

Let us now review once more the various conclusions of Li and Zaidi within the context of our present demonstrations. The evidence is quite clear that observers can perceive 3D shape from texture under both perspective and orthographic projections. Perspective projections are only necessary in the highly unusual circumstance where surface contours are all aligned in planes that are parallel to the line of sight. If a curved surface with gradual orientation changes is viewed from a non-degenerate orientation (see Fig. 11), if its contours are non-planar (see Fig. 10), or if it is textured with random polka dots (see Fig. 7), then the appearance of 3D shape from texture can be quite similar for both perspective and orthographic projections.

The evidence is also clear that the perception of 3D shape from texture does not necessarily require discrete oriented energy in the directions of principal curvature or a frequency spectrum that is elongated in the critical directions. For example, the diagonal contour textures in Fig. 10 have no energy at all in the critical directions, and the polka dot textures in Figs. 7 and 15 have uniform orientation spectra, yet both of these textures provide perceptually compelling information about the 3D surface structure.

Not only is oriented energy in the direction of principal curvature not necessary for the perceptual appearance of 3D form, it is also unnecessary to obtain a stable and accurate perception of the direction of surface slant or the sign of surface curvature. Li and Zaidi's analysis ignores several other potential sources of information about these properties that have been identified previously in the literature. These include gradients of contour compression or convergence (Attneave & Olson, 1966; Gillam, 1968, 1970), gradients of texture size or shape (Knill, 1998a,b,c; Stevens, 1981b), and the presence of smooth occlusion contours (Koenderink & van Doorn, 1982; Koenderink, 1984). Other research has shown, in addition, that the perceptual mechanisms for determining surface slant or the sign of surface curvature can be strongly influenced by perceptual biases to prefer convex surfaces over concave surfaces and ground surfaces over ceiling surfaces. These biases reflect the relative frequency of occurrence of these different surface types in the natural environment. When they are in conflict with other sources of texture information the apparent 3D shape of a surface can become perceptually multi-stable.

The human visual system is remarkably robust in its ability to analyze patterns of optical texture to determine the 3D structures of visible surfaces. The perception of 3D shape from texture can occur for 2D mapped textures, such as wall paper (e.g. Li & Zaidi, 2000), or 3D volume textures, such as wood or marble (e.g. Cumming, et al., 1993), and those textures can be composed of random blobs, such as polka dots (Knill, 1998a,b,c; Phillips, 1970; Rosenholtz & Malik, 1997), or patterns of continuous contours (Attneave & Olson, 1966; Stevens, 1981a; Todd & Reichel, 1990). Observers can exploit texture information to perceive the structure of planar surfaces (e.g. Gibson, 1950b; Knill, 1998a,b,c; Rosenholtz & Malik, 1997), surfaces that are singly curved, such as cylinders (e.g. Cumming et al., 1993; Cutting & Millard, 1984; Knill, 2001), or surfaces that are doubly curved, such as ellipsoids (Todd & Akerstrom, 1987; Todd & Reichel, 1989, 1990). Finally, in the case of curved surfaces, the perception of shape from texture can occur for images produced with either perspective or orthographic projections (Knill, 2001; Todd & Akerstrom, 1987; Todd & Reichel, 1989, 1990).

In contrast to the robustness with which human observers are able to perceive shape from texture, most computational analyses of this process can only be used for a limited class of surfaces and/or a limited class of textures. The earliest techniques for computing shape from texture (e.g. Purdy, 1958; Witkin, 1981) were based on an assumption that the texture is isotropic – that is to say, that all texture elements are approximately circular. These models do not work effectively with anisotropic contour textures like the ones shown in Figs. 8–12. Other analyses that have been designed specifically for contour textures have been based on an assumption that the depicted surface is singly curved, and that the contours lie along lines of curvature (Stevens, 1981a) or surface geodesics (Knill, 2001). These models do not work effectively with doubly curved surfaces like the one in Fig. 12, or with surfaces that have polka dot textures like the one in Fig. 7. A more promising approach, we believe, is the one developed by Malik and Rosenholtz (1997), which computes local surface structure by measuring the affine distortions between texture patches in neighboring image regions (see also Gårding, 1992, 1993), based on a more ecologically reliable assumption about texture homogeniety.

Of all the images we have presented in this discussion, the one shown in the left panel of Fig. 12 poses the greatest difficulty for current theoretical approaches, because it violates almost all of the assumptions that have previously been employed for the computational analysis of shape from texture. This image depicts a doubly curved surface under orthographic projection that is textured with an anisotropic contour pattern. The contours do not lie along lines of curvature or surface geodesics, and they are not homogeneously distributed over the surface. One possible strategy for analyzing this type of display might be to compute 3D structure based on an assumption that the contours are planar cuts (see Tse, 2001). The problem with this approach, however, is that it would only work for a limited class of textures, and would not be appropriate for the helical contour patterns in Fig. 10 or the polka dot patterns in Figs. 7 and 15. How human observers are able to perceive shape from texture over such a wide variety of conditions remains as an important problem for future research.

Acknowledgements

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

References

Attneave, F., & Olson, R. K. (1966). Inferences about visual mechanisms from monocular depth effects. *Psychnomic Science*, *4*, 246–256.

Braunstein, M. L. (1976). *Depth perception through motion*. New York: Academic Press.

Buckley, D., & Frisby, J. P. (1993). Interaction of stereo, texture, and outline cues in the shape perception of three-dimensional ridges. *Vision Research*, *33*, 919–934.

Cumming, B. G., Johnston, E. B., & Parker, A. J. (1993). Effects of different texture cues on curved surfaces viewed stereoscopically. *Vision Research*, *33*, 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*, 198–216.

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 the visual world*. Boston, MA: Haughton Mifflin.

Gibson, J. J. (1950b). The perception of visual surfaces. *American Journal of Psychology*, *63*, 367–384.

Gillam, B. (1968). Perception of slant when perspective and stereosis conflict: experiments with aniseikonic lenses. *Journal of Experimental Psychology*, *78*, 299–305.

Gillam, B. (1970). Judgments of slant on the basis of foreshortening. *Scandinavin Journal of Psychology*, *11*, 31–34.

Knill, D. C. (1998a). Surface orientation from texture: Ideal observers, generic observers and the information content of texture cues. *Vision Research*, *38*, 1655–1682.

Knill, D. C. (1998b). Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Research*, *38*, 1683–1711.

Knill, D. C. (1998c). 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.

Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, *30*, 321–330.

Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, *2*, 129–137.

Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1996). Pictorial surface attitude and local depth comparisons. *Perception and Psychophysics*, *58*, 163–173.

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*, 431–448.

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.

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. (1998). Observer biases in the 3D interpretation of line drawings. *Vision Research*, *38*, 2817–2832.

Phillips, R. J. (1970). Stationary visual texture and the estimation of slant angle. *Quarterly Journal of Experimental Psychology*, *22*, 389–397.

Purdy, W. C. (1958). The hypothesis of psychophysical correspondence in space perception. Unpublished Ph.D., Cornell University.

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.

Reichel, F. D., Todd, J. T., & Yilmaz, E. (1995). Visual discrimination of local surface depth and orientation. *Perception and Psychophysics*, *57*, 1233–1240.

Rosenholtz, R., & Malik, J. (1997). Surface orientation from texture: isotropy or homogeneity (or both)? *Vision Research*, *37*, 2283–2293.

Stevens, K. A. (1981a). The visual interpretation of surface contours. *Artificial Intelligence*, *17*, 47–53.

Stevens, K. A. (1981b). The information content of texture gradients. *Biological Cybernetics*, *42*, 95–105.

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., & 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. (2001). A planar cut approach to volume recovery from silhouettes. *Psychological Review*, in press.

Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence*, *17*, 17–45.