



# Information limitations in perception of shape from texture

Andrea Li \*, Qasim Zaidi

*SUNY State College of Optometry, 33 W 42th Street, New York, NY 10036-8003, USA*

Received 26 June 2000; received in revised form 9 January 2001

---

## Abstract

Li and Zaidi (Li, A., and Zaidi, Q. (2000) *Vision Research*, 40, 217–242) showed that the veridical perception of the 3-dimensional (3D) shape of a corrugated surface from texture cues is entirely dependent on the visibility of critical patterns of oriented energy. These patterns are created by perspective projection of surface markings oriented along lines of maximum 3D curvature. In images missing these orientation modulations, observers confused concavities with convexities, and leftward slants with rightward slants. In this paper, it is shown that these results were a direct consequence of the physical information conveyed by different oriented components of the texture pattern. For texture patterns consisting of single gratings of arbitrary spatial frequency and orientation, equations are derived from perspective geometry that describe the local spatial frequency and orientation for any slant at any height above and below eye level. The analysis shows that only gratings oriented within a few degrees of the axis of maximum curvature exhibit distinct patterns of orientation modulations for convex, concave, and leftward and rightward slanted portions of a corrugated surface. All other gratings exhibit patterns of frequency and orientation modulations that are distinct for curvatures on the one hand and slants on the other, but that are nearly identical for curvatures of different sign, and nearly identical for slants of different direction. The perceived shape of surfaces was measured in a 5AFC paradigm (concave, convex, leftward slant, rightward slant, and flat-frontoparallel). Observers perceived all five shapes correctly only for gratings oriented within a few degrees of the axis of maximum curvature. For all other oriented gratings, observers could distinguish curvatures from slants, but could not distinguish signs of curvature or directions of slant. These results demonstrate that human observers utilize the shape information provided by texture components along both critical and non-critical orientations. © 2001 Elsevier Science Ltd. All rights reserved.

*Keywords:* Shape from texture; Axis of maximum curvature; Orientation modulations; Spatial frequency modulations

---

## 1. Background

When a 3-dimensional (3D) surface is projected in perspective, variations in the surface pattern in the 2-dimensional (2D) image provide potential cues to the shape of the surface (Gibson, 1950). Li and Zaidi (2000), however, discovered that observers perceived the correct 3D shape of the surface only for some surface textures. Concavities and convexities were correctly distinguished only when the image contained patterns of orientation modulations corresponding to projected lines of maximum surface curvature, whereas frequency modulations were insufficient. By characterizing shape information in terms of local frequency and orientation, features which primary visual cortex is

well-equipped to extract, these results provided empirical grounds for a neural model of extracting 3D percepts from 2D texture cues.

In this paper, it will be formally shown 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. All other oriented texture components form distinct patterns for curvatures, and planar slanted portions, but nearly identical patterns for curvatures of different sign and nearly identical patterns for slants of different direction. For these components, observers use the information to distinguish between curvatures and slants, but are unable to distinguish signs of curvatures or directions of slants. For compound textures, observers combine information from texture components

---

\* Corresponding author. Fax: +1-212-7805137.

E-mail address: ali@sunyopt.edu (A. Li).

along critical and non-critical orientations in extracting the 3D shape of the surface.

## 2. Orientation and frequency modulations

Li and Zaidi (2000) used a surface that was corrugated in depth as a sinusoidal function of the horizontal axis. This surface, like flat and cylindrical surfaces, is *developable*, i.e. curved along, at most, one dimension in 3D space. Through every point on this surface passes a line of zero curvature, and orthogonal to it, a line of maximum curvature. Developable surfaces can be formed by bending a piece of paper without stretching, tearing, or crumpling, thus leaving the texture pattern on the surface unchanged. Except for a few studies (Stevens, 1981; Reichel & Todd, 1990; Todd & Reichel, 1990; Mamassian & Landy, 1998), previous shape-from-texture studies have utilized flat or singly curved surfaces such as cylinders (Braunstein & Payne, 1969; Gillam, 1970; Vickers, 1971; Bajcsy & Lieberman, 1976; Rosinki & Levine, 1976; Cutting & Millard, 1984; Todd & Akerstrom, 1987; Turner, Gerstein, & Bajcsy, 1991; Blake, Bulthoff, & Sheinberg, 1993; Cummings, Johnston, & Parker, 1993; Sakai & Finkel, 1993; Krumm & Shafer, 1994; Malik & Rosenholtz, 1997; Goodenough & Gillam, 1997; Knill, 1998a,b,c). However, critical tests of shape-from-texture theories require a surface that contains both concavities and convexities.

Fig. 1A and B are perspective images of a corrugated surface textured with a horizontal–vertical plaid pattern. In Fig. 1A, the simulated surface contains a concavity along the central vertical mid-line of the image, and in Fig. 1B the simulated surface contains a central convexity. Both images exhibit orientation modulations of the horizontal component of the plaid and frequency modulations of the vertical component. The horizontal component is oriented parallel to the axis of maximum curvature and the vertical component parallel to the axis of zero curvature. The horizontal components are shown in isolation in Fig. 1C and D, and the vertical components in Fig. 1E and F. When presented with just the orientation modulations, observers perceived convexities and concavities at the correct locations along the surface. When presented with just the frequency modulations, observers correctly perceived the locations of curvatures, but confused concavities and convexities (Li & Zaidi, 2000). Although the frequency modulations cannot be used to differentiate between concavities and convexities, in Fig. 1A and B when they are added to the orientation modulations, the amplitude of the depth corrugation appears increased and the shape appears sharper compared to Fig. 1C and D.

In Fig. 1G, the orientation modulations from the projected concavity (Fig. 1C) have been added to the

frequency modulations from the projected convexity (Fig. 1F), and in Fig. 1H, the frequency modulations from the projected concavity (Fig. 1E) have been added to the orientation modulations from the projected convexity (Fig. 1D). In both Fig. 1G and H, the signs of the perceived curvatures are entirely dictated by the orientation modulations; Fig. 1G appears to contain a central concavity while Fig. 1H appears to contain a central convexity. While there are small differences between Fig. 1A and G and between Fig. 1B and H, the frequency modulations in all four images only serve to sharpen the perceived depth so that the surface appears more triangular compared to the corrugated horizontal components.

These results apply not only to simple ruled patterns, but also to more complex texture patterns. The ‘octotropic plaid’ pattern in Fig. 2A and B is composed of eight compound gratings, spaced equally apart in orientation, each component consisting of three frequencies at random phases. Before corrugation, the octotropic plaid contains symmetric rosette-like patterns that are unstable, similar to Marroquin patterns (Marroquin, 1976). When the horizontal component of the octotropic plaid was presented in isolation (Fig. 2C and D) observers perceived concavities and convexities at the correct locations. For the remaining seven components in Fig. 2E and F, which contain all the relevant texture gradients and frequency modulations consistent respectively with a central concavity and convexity, observers reported both central curvatures as convex or both as concave. Subtracting any component other than the horizontal, resulted in shape percepts similar to those in Fig. 2E and F (Li & Zaidi, 2000). When the proper patterns of both orientation and frequency modulations are present (Fig. 2A and B), the corrugation appears sharper and increased in amplitude relative to Fig. 2C and D. In Fig. 2G and H, the horizontal component and the remaining 7 components consistent with opposite phases of the corrugation have been added together. The signs of the perceived curvatures of each image are dictated by the orientation modulations of the horizontal component; Fig. 2G appears to contain a central concavity while Fig. 2H appears to contain a central convexity.

## 3. Componential analysis

The goal of this paper is to quantify all the shape information contained in different oriented components of a texture pattern. We will thus derive the orientation and frequency modulations in texture components oriented along axes of principal (i.e. maximum and zero) curvature, and non-principal curvature. In a psychophysical experiment, we will measure how these patterns of modulations contribute to the overall 3D percept of the surface.

The octotropic plaid pattern is ideal for these purposes because it is a complex texture that consists of discrete grating components. In this section, we derive equations for the local spatial frequency and orientation in the projected image, given the orientation of

the texture component and the local slant of the surface. The equations hold for perspective projections of any developable surface in the fronto-parallel position (i.e. for which tilt = 0 at all points on the surface).




Fig. 1. (a) A horizontal-vertical plaid pattern drawn on a flat surface which was then corrugated in depth as a sinusoidal function of horizontal position, and projected in perspective into the image plane. Observers correctly identify the central concavity in the surface. (b) The same textured surface corrugated with a central convexity. Observers correctly identify the central convexity. (c–d) The horizontal grating components of the plaids in a and b, respectively. Contours are formed by orientation modulations following projected lines of maximum curvature. Observers correctly identify the central concavity in c and convexity in d. (e–f) The vertical grating components of the plaids in a and b, respectively. Contours along projected lines of zero curvature form frequency modulations. Observers incorrectly report a central convexity in both images. (g) Orientation modulations from the projected concavity in c added to frequency modulations from the projected convexity in f. The perceived curvature is dictated by the orientation modulations. (h) Orientation modulations from the projected convexity in d added to frequency modulations from the projected concavity in e. The perceived curvature is dictated by the orientation modulations.




Fig. 2. (a) An 'octotropic plaid' pattern corrugated and projected as the horizontal-vertical plaid pattern in Fig. 1a. The pre-corrugated pattern is the sum of eight components, each oriented  $22.5^\circ$  from the next, where each component is a complex grating composed of three frequencies. Observers correctly identify the central concavity in the surface. (b) The same pattern corrugated with a central convexity. Observers correctly identify the central convexity. (c–d) The horizontal components of the plaids in a and b, respectively. Contours are formed by orientation modulations following projected lines of maximum curvature. Observers correctly identify the central concavity in c and convexity in d. (e–f) The remaining seven components of the plaids in a and b, respectively. Observers incorrectly report a central convexity in both images. (g) The horizontal component from the projected concavity in c added to the remaining seven components from the projected convexity in f. The perceived curvature is dictated by the orientation modulations of the horizontal component. (h) The horizontal component from the projected convexity in d added to the remaining seven components from the projected concavity in e. The perceived curvature is dictated by the orientation modulations of the horizontal component.

### 3.1. Derivation of projected local orientation and frequency

The reader is referred to Fig. 3 for the geometry of this derivation. The eye is located at  $P$ . The image

plane is at distance  $D$ . A texture component at angle  $\omega$  with respect to the horizontal falling on a surface slanted at angle  $\theta$  from the fronto-parallel plane was locally approximated by a line  $H$  of unit length on the vertical axis through the center of the image beginning

at a height  $y$  with respect to eye height, oriented in the  $xy$ -plane at  $\omega$  and slanted in the  $xz$ -plane at  $\theta$ . The line  $H$  projects in the image to  $h$  which is oriented at angle  $\eta$  with respect to the horizontal:

$$\eta = \tan^{-1} \left( \frac{D \sin \omega - y \cos \omega \sin \theta}{D \cos \theta \cos \omega} \right) \quad (1)$$

The length of  $h$  is proportional to the inverse of the projected local spatial frequency of the grating component *perpendicular* to  $\omega$ . Taking this into account, the local projected frequency ( $F'$ ) of a component oriented at  $\omega$  is derived as:

$$F' = \frac{F(D + \cos \omega_p \cos \theta)}{\sqrt{\cos^2 \theta (\cos^2 \omega_p (D^2 + y^2)) + D^2 \sin^2 \omega_p - 2yD \sin \omega_p \cos \omega_p \sin \theta}} \quad (2)$$

where  $F$  is the frequency of the uncorrugated component and  $\omega_p$  is  $\omega + \pi/2$ .

### 3.2. Illustrations of componential analysis

To illustrate the derivations, we used a surface that was corrugated in depth as a sinusoidal function of the horizontal axis ( $x$ ). The distance ( $d$ ) between the observer and the surface along the observer's line of sight was computed as:

$$d = A \cos(2\pi fx + \varphi) + D \quad (3)$$

where  $A$  is the amplitude of the depth modulation and is set to 1,  $f$  is the frequency (0.58 cycles per horizontal extent of the image),  $\varphi$  is the phase of the corrugation at the center of the image, and  $D$  is distance between the eye and the image plane. The amplitude of the depth modulation was computed to be 8 cm from peak to trough for a viewing distance of 44 cm. Texture patterns on this surface consisted of sinusoidal gratings at each of five equally spaced orientations from horizontal to vertical (0, 22.5, 45, 67.5, and 90°). The frequency of the gratings before corrugation was 0.85 cpd, and phase was randomized. The surface was projected in perspective onto the image plane of a CRT monitor (see Appendix A). The nodal point of the eye was defined as the origin. The center of each image was aligned in height and position with the observer's eye. Each image was  $79 \times 381$  pixels and subtended  $4.3^\circ \times 21^\circ$  of visual angle. The surface was corrugated at one of four different central phases: a concavity ( $\varphi = 0$ ), a convexity ( $\varphi = \pi$ ), a leftward slant ( $\varphi = \pi/2$ ) and a rightward slant ( $\varphi = 3\pi/2$ ). Fig. 4 illustrates how the four surface shapes were projected with respect to the observer. Projections were computed using Matlab.




Fig. 3. Geometry used in the derivation of projected orientation ( $\eta$ ) of line  $H$ , tilted in the image plane at the same orientation as the pre-corrugated component ( $\omega$ ), and slanted in the  $xz$ -plane at the same angle as the local slant of the corrugated surface ( $\theta$ ).




Fig. 4. The four half-cycle portions of the corrugated surface used in the experiments. See text for details.

#### 3.2.1. Components parallel to axes of principal curvature

Projections of the four different surface shapes for the horizontal (0°) component are shown in the top row of Fig. 5. The curved contours in the images are formed by local changes in orientation, while the spatial frequency or the width of the individual contours remains nearly constant. The pattern of orientation modulations is distinct for each of the four surface shapes.

The curvature of each contour in these four images can be characterized as level, bowed upwards, bowed downwards, positively oblique (where horizontal is defined at 0°), or negatively oblique. These forms are projections of different 3D surface shapes, depending on position with respect to the level of the observer's




Fig. 5. Projected images of the four shapes from Fig. 4 for components oriented at 0 (horizontal), 22.5, 45, and 67.5, and 90° (vertical). Each image of the 0° component exhibits a distinct pattern of orientation modulations following projected lines of maximum surface curvature. Images of the four shapes for the 22.5, 45, and 67.5° components exhibit concurrent orientation and frequency modulations the patterns of which are similar for both curvatures and for both slants. Images of the four shapes for the 90° component exhibit nearly identical patterns of frequency modulations for both curvatures and for both slants.

eye. Fig. 6a shows a horizontal planar cross section below the observer's eye level containing a concavity (arc BAC) and a convexity (arc BDC) at the same horizontal position. Under perspective projection, parallel lines L1, L2, and L3 project to lines converging at the center of the image,  $p$ . By connecting each of the points A, B, C, and D to the eye through the projections of lines L1, L2, and L3, it is seen that the concavity projects to an upwardly bowed contour (bac) and the convexity projects to a downwardly bowed contour (bdc), consistent with the farther point projecting closer to the projection of the line of sight at eye height (see Mundy & Zisserman, 1992). Contours on the surface following the concavity and the convexity therefore project to two distinctly bowed contours in

the image. Turning this figure upside down shows that the bowing of the projected contours in the image *above* the observer's eye level reverses, i.e. the concavity projects to a downwardly bowed contour, while the convexity projects to an upwardly bowed contour. Similarly in Fig. 6b, a rightward slant (EF) projects to a positively oblique line (ef), and a leftward slant (GH) projects to a negatively oblique line (gh). Contours on the surface along leftward and rightward slants thus project to distinctly oriented lines in the image.

The top and bottom panels of Fig. 7 show respectively the local orientations ( $\eta$ ) and spatial frequencies ( $F'$ ), derived in Eqs. (1) and (2), plotted as a linear function of surface slant ( $\theta$ ) at heights of  $y = +10^\circ$  (axis along top of the graphs) and  $y = -10^\circ$  (axis along bottom of the graphs) in the image. Different symbols represent values for the five different grating components. (Angles of image height (in the  $yz$ -plane) and angles of surface slant (in the  $xz$ -plane) will be expressed in terms of 'deg'; the orientation of the uncorrugated components and local orientation values in the image (both in the  $xy$ -plane) will be expressed with the symbol '°'.) Horizontal lines indicate the orientation of each of the five uncorrugated components. Local projected orientation values for slants spanning a concavity are plotted in the top left panel, and those for slants spanning a convexity in the top right panel. A leftward slant spans a slant of  $0^\circ$  in the left panel to  $0^\circ$  in the right panel, and a rightward slant spans  $0^\circ$  in the right panel to  $0^\circ$  in the left panel (looping around). The value of  $\eta$  was calculated for projected lines originating on the vertical mid-line of the image at a distance  $D$ . The actual values of  $\eta$  will differ slightly from this calculation for locations far from the vertical mid-line due to perspective and distance changes. However, within the  $\pm 2^\circ$  span of our images, the differences are negligible.

Values for the projected horizontal ( $0^\circ$ ) component are plotted in solid triangles. The patterns of orientation modulations for this component are distinct for the concavity and for the convexity. The local orientation values at the steepest slants (slant =  $+84^\circ$  and  $-84^\circ$ ) are equal in magnitude but opposite in sign ( $-59^\circ$  and  $+59^\circ$ ). The local orientation at fronto-parallel portions of the surface (slant =  $0^\circ$ ) equals the orientation of the uncorrugated component, i.e.  $0^\circ$ . The values across the top of the concavity ( $y = +10^\circ$ ) thus change from negative to positive, and those across the convexity change from positive to negative. The orientation values across each of the two slants are either all positive (leftward slant) or all negative (rightward slant), resulting in a distinct pattern for each slant. As illustrated in Fig. 6, in perspective projection, the local orientations below the observer's line of sight are equal in magnitude but opposite in sign from those above the observer's line of sight. Thus along the bottom of the image

( $\gamma = -10^\circ$ ), the signs of the projected orientation values are reversed in sign from those along the top of the image. In summary, the 2D patterns of orientation modulations are distinct for each of the four surface shapes.

Local spatial frequency is similarly plotted as a function of surface slant in the bottom panels of Fig. 7. At both image heights, values across the projected horizontal component are fixed at 0.85 cpd, the frequency of the uncorrugated component. (Small deviations due to perspective are smaller than the size of the symbols.)

The projected shapes for the vertical ( $90^\circ$ ) component are shown in the bottom row of Fig. 5. Unlike the projected horizontal component for which the images of the four surface shapes were distinct, the projected vertical component yields nearly identical images for the concavity and convexity, and nearly identical images for the two slants. Fig. 7 (solid circles) shows that the pattern changes across these images consist entirely of frequency modulations, while the local orientation at both heights is fixed at  $90^\circ$ . The bottom panels of Fig. 7 show that the patterns of frequency modulations are nearly identical for the concavity and the convexity (changing from high to low to high), and nearly identical for the two slants (changing from low to high to low), but distinct for curvatures as compared to slants. In the bottom row of Fig. 5, though there are small differences in frequency due to slight differences in distance between the observer and the surface, differences between the images of the two curvatures and between the images of the two slants are almost entirely due to the different phases of the pattern on the surfaces.

We now examine the patterns of orientation and frequency modulations exhibited by components oriented along lines of non-principal surface curvature.

### 3.2.2. Components parallel to axes of non-principal curvature

The projected shapes for the  $22.5^\circ$  component are shown in the second row of Fig. 5. The derived orientation and frequency values are plotted as open circles in Fig. 7. The orientation values at the most slanted portions of the surface (slant =  $+84^\circ$  and  $-84^\circ$ ) are unequal but of the same sign, and steeper than the uncorrugated orientation ( $67^\circ$  and  $80^\circ$ ). Despite the asymmetry, the images of the two curvatures in Fig. 5 exhibit similar patterns of curved contours with nearly parallel flow patterns across the vertical extent of each image. A different pattern of contours is exhibited for images of the two slants, but the pattern is similar across the two. For all four surface shapes, the contours are closest to horizontal along the central vertical strip of the concavity and the convexity, and along the left and right edges of the rightward and leftward slant. The patterns of orientation modulations for the concavity and the convexity are thus nearly identical, increasing almost symmetrically from  $22.5^\circ$  at the center of each image, out towards the left and right edges. The pattern of orientation modulations is different for the slants than for the curvatures, but is nearly identical for the two slant directions, decreasing from  $67^\circ$  or  $80^\circ$  at the center out towards the edges of the image. The bottom panels of Fig. 7 show that there are small frequency modulations (between 0.85 and 1.0 cpd) that are asymmetric at leftward and rightward slants, but they do not result in qualitatively different patterns between the two curvatures or between the two slants.

The projected shapes for the  $45^\circ$  and  $67.5^\circ$  components are shown in the third and fourth rows of Fig. 5. For each component, images of the concavity and convexity are nearly identical, as are those of the two




Fig. 6. (a) Planar cross-section, below the observer's eye level, through a concavity (BAC) and a convexity (BDC). The perspective projections of lines L1, L2, and L3 converge at point p at the center of the image. The concavity (BAC) projects to an upwardly-bowed contour (bac) in the image, and the convexity (BDC) projects to a downwardly-bowed contour (bdc). Turning this figure upside-down shows that the bowing of the projected contours with respect to the observer reverses for cross-sections above the observer's eye level. Concavities and convexities thus project to two distinct patterns of contours in the image. (b) Similarly, slants of different direction project to distinct patterns of oblique contours in the image.

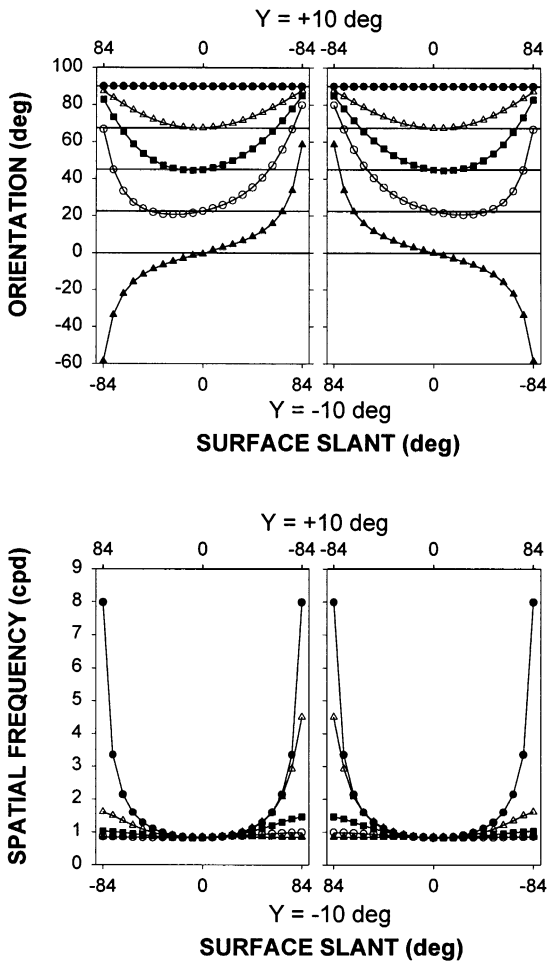


Fig. 7. *Top*: Derived local orientation values in the image as a linear function of the slant of the surface for each of the five oriented components, across image heights of  $+10$  (top abscissa) and  $-10^\circ$  (bottom abscissa) from eye level in the image. Derived orientation values for each of the five oriented components are plotted in different symbols:  $0^\circ$  in solid triangles,  $22.5^\circ$  in open circles,  $45^\circ$  in solid squares,  $67.5^\circ$  in open triangles, and  $90^\circ$  in solid circles. Values across the projected concavity are plotted in the left panel, and values across a convexity are plotted in the right panel. Values across the leftward slant fall between a slant of  $0^\circ$  in the left panel to a slant of  $0^\circ$  in the right panel, and values across the rightward slant fall between a slant of  $0^\circ$  in the right panel to a slant of  $0^\circ$  in the left panel (wrapping around). At fronto-parallel portions of the surface (at slants of  $0^\circ$ ), the derived orientation equals the orientation of the grating in uncorrugated form, indicated by the solid horizontal lines. Only the projected horizontal component (solid triangles) exhibits local orientations of opposite sign at leftward and rightward slants (i.e. slants of  $-84$  and  $+84^\circ$ , respectively). All other components exhibit nearly identical orientations at slants of  $+84$  and  $-84^\circ$  that are of the same sign and are steeper than the uncorrugated orientation. *Bottom*: Derived local spatial frequency in the images as a function of the slant over the same range of slants and image heights as in the top panels. At fronto-parallel portions of the surface (at slants of  $0^\circ$ ), the frequency for all components equals the uncorrugated frequency ( $0.85$  cpd). Frequencies increase as the slant of the surface deviates from fronto-parallel.

slants. Fig. 7 (solid squares and open triangles) shows that, like the  $22.5^\circ$  component, the local orientations at

slants of  $+84$  and  $-84^\circ$  are of the same sign and steeper than the uncorrugated orientation and are nearly equal. The patterns of orientation modulations are thus nearly identical for the concavity and the convexity, and for the two slants, but different for curvatures as compared to slants. The bottom panels of Fig. 7 show greater frequency modulations compared to the  $22.5^\circ$  component. Despite the small asymmetries between the frequencies at slants of  $+84$  and  $-84^\circ$ , the images of the two curvatures are highly similar, as are those of the two slants.

#### 4. Experiment 1: Global shape percepts

##### 4.1. Psychophysical method

The contribution of each oriented component to the 3D percept of the surface was measured using a global shape identification task. All the images in Fig. 5 were presented plus images of each component corresponding to a flat fronto-parallel surface. Images were presented on a SONY GDM-F500 flat screen monitor with a  $800 \times 600$  pixel screen running at a refresh rate of 100 frames/second via a Cambridge Research Systems Video Stimulus Generator (CRS VSG2/3) controlled through a 400 MHz Pentium II PC. Through the use of 12-bit DACs, after gamma correction, the VSG2/3 is able to generate 2861 linear levels for each gun. The mean luminance of the screen was  $22.4$   $\text{cd}/\text{m}^2$ .

Observers identified the perceived shape of the surface in a 5AFC task (convexity, concavity, leftward slant, rightward slant, and flat-frontoparallel). A total of 75 conditions (5 component patterns  $\times$  5 shapes  $\times$  3 carrier phases) at ten trials each were randomly interleaved for a total of 750 trials. These were divided into two separate sessions, each lasting approximately 20 min. At the beginning of each session, observers adapted for 1 min to a uniform DC level strip of the same dimensions as the stimulus, containing a central fixation. Each image was then presented with a central fixation for 1 s, after which the screen returned to the uniform strip until a response was made. Using a three-switch response box, observers indicated whether they saw a concavity, a convexity, leftward slant, rightward slant, or a flat-frontoparallel surface. Each switch on the box could be flipped either towards or away from the observer, so response configurations were chosen to match the different shapes with respect to the observer: a concavity was indicated by pulling the left and right switches toward the observer, a convexity by pushing them away, a rightward slant by pulling the left switch towards and pushing the right switch away, and vice versa for a leftward slant; a flat-frontoparallel surface was indicated by pulling or pushing the middle switch. Observers were instructed to choose the shape



that best resembled their percept. Viewing was monocular, and the head was stabilized with a chinrest.

The observers consisted of one of the authors (AL), and two observers (FG and JM), both experienced in psychophysical tasks but naive about the purposes of the experiment. All had corrected-to-normal acuity.

#### 4.2. Results

Data for the three observers are shown in Fig. 8. In each panel, perceived shape is plotted against simulated shape. The area of each disk represents the frequency with which each perceived shape is reported, with the area totaling unity along each simulated shape category. If all shapes are identified correctly, the plot will show large disks confined to the diagonal; incorrect identifications will be represented by disks off of the diagonal.

Data for the horizontal component (Fig. 5, top row) are shown in Fig. 8a. For this component, all three observers identified the five shapes correctly.

Fig. 8e shows data for the vertical component in (Fig. 5, bottom row). Observer JM identified all shapes as flat-frontoparallel, and was thus unable to distinguish curvatures from slants. Observer FG predominantly identified both concave and convex curvatures as convex, and both rightward and leftward slants as flat-frontoparallel. He was thus able to distinguish curvatures from slants, but unable to distinguish the signs of the curvatures or the directions of the slants. Observer AL was also able to distinguish curvatures from slants, but was unable to distinguish the signs of the two curvatures, identifying them both as convex. The two slants were also not distinguished, and were identified as concave.

Data for the 22.5° component (Fig. 5, second row) are shown in Fig. 8b. Observers were unable to distinguish between signs of curvatures or directions of slants. Observer JM identified curvatures as either convex or rightward slanted, and both slants as rightward slanted. Observer FG identified both curvatures and both slants as either leftward or rightward slanted. Observer AL identified both curvatures as convex, and both slants as rightward slants.

Data for the 45° component (Fig. 5, third row) are shown in Fig. 8c. All observers were unable to distinguish between signs of curvatures and directions of slants. Observer JM identified both curvatures as convex, and both slants as slanted either leftward or rightward, but could not consistently distinguish between the two. Observer FG was able to distinguish between curvatures and slants, but both curvatures were identified as rightward slants, and both slants were identified as leftward slants. Observer AL identified both curvatures as convex and both slants as concave curvatures.

Data for the 67.5° component (Fig. 5, fourth row) are shown in Fig. 8d. Observer JM was able to distinguish curvatures from slants, but was unable to distinguish between the different curvatures and the different slants. Both curvatures were identified as convex and both slants were identified as flat-frontoparallel. While observer FG identified some portion of the trials for each shape correctly, he was unable to consistently distinguish concave from convex curvatures and the two different slants. Observer AL identified both curvatures as convex and both slants as concave curvatures.

Taken together, these data show that observers use the shape information conveyed by both critical and non-critical components. Eq. (1) and the derived plots in Fig. 7 show that patterns of orientation modulation of the horizontal component are distinct for each of the five shapes. The results for the horizontal component in Fig. 8a show that observers use these differences not only to distinguish, but to correctly identify each shape. The derivation of patterns of orientation and frequency modulations of all other components in Fig. 7 show qualitative differences between curvatures and slants, but are nearly identical for concavities and convexities on the one hand, and leftward and rightward slants on the other. Fig. 8b–e shows that observers utilize this information, at best, to distinguish curvatures from slants, which is all that is supported by this information. For all components other than the horizontal one, observers could not distinguish signs of curvatures or directions of slants, and in some cases curvatures and slants were completely indistinguishable (e.g. for the 22.5° component for observer FG, or the vertical component for observer JM).

#### 5. Experiment 2: Components close to the axis of maximum curvature

How far from the axis of maximum curvature can the orientation of a component deviate before it stops conveying information about the signs of curvatures and slants? This will depend on the exact shape of the surface and the distance between the surface and the observer. To illustrate the nature of this dependence, we examined two sinusoidal surfaces with a corrugation frequency of 0.14 cycles per degree calculated for a viewing distance of 107 cm: the first with an amplitude of one, and the second with an amplitude of two. Projected images of the latter surface were optically similar to those of the surface in Experiment 1 (Fig. 5). Texture patterns consisted of sinusoidal gratings of 2 cpd oriented at 0, 3, 6, 9, and 12°. The relationships between surface slant and image position for a sinusoidal surface with an amplitude of one are plotted as the thick solid lines in the top panels of Fig. 9: concavity on the left and convexity on the right. Image width

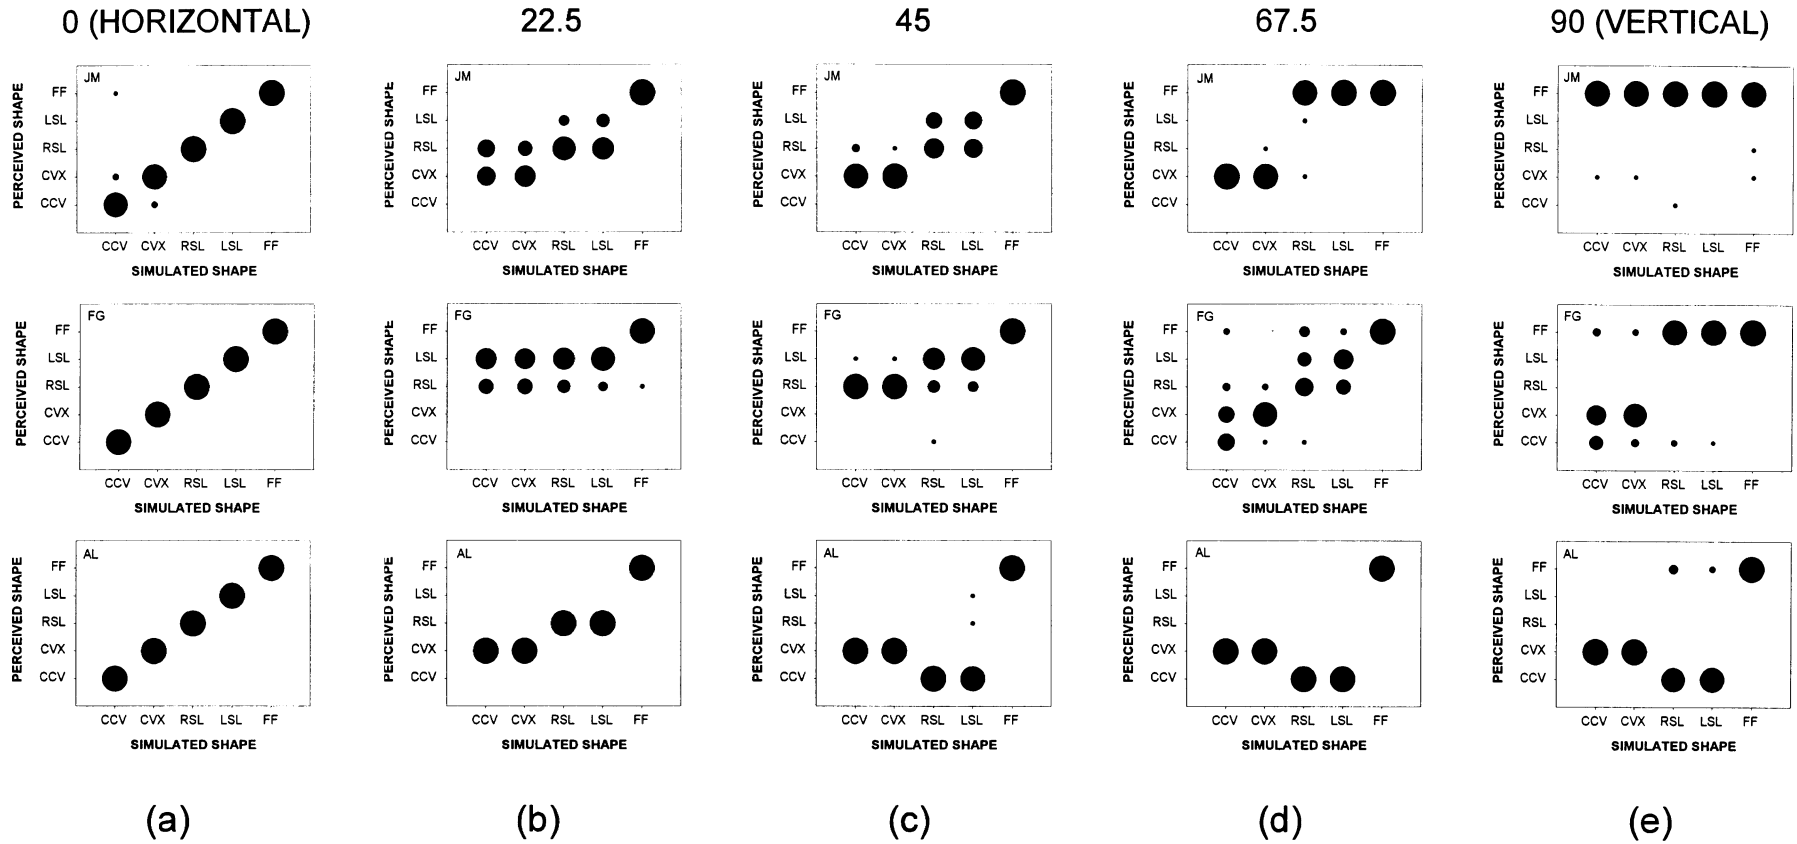
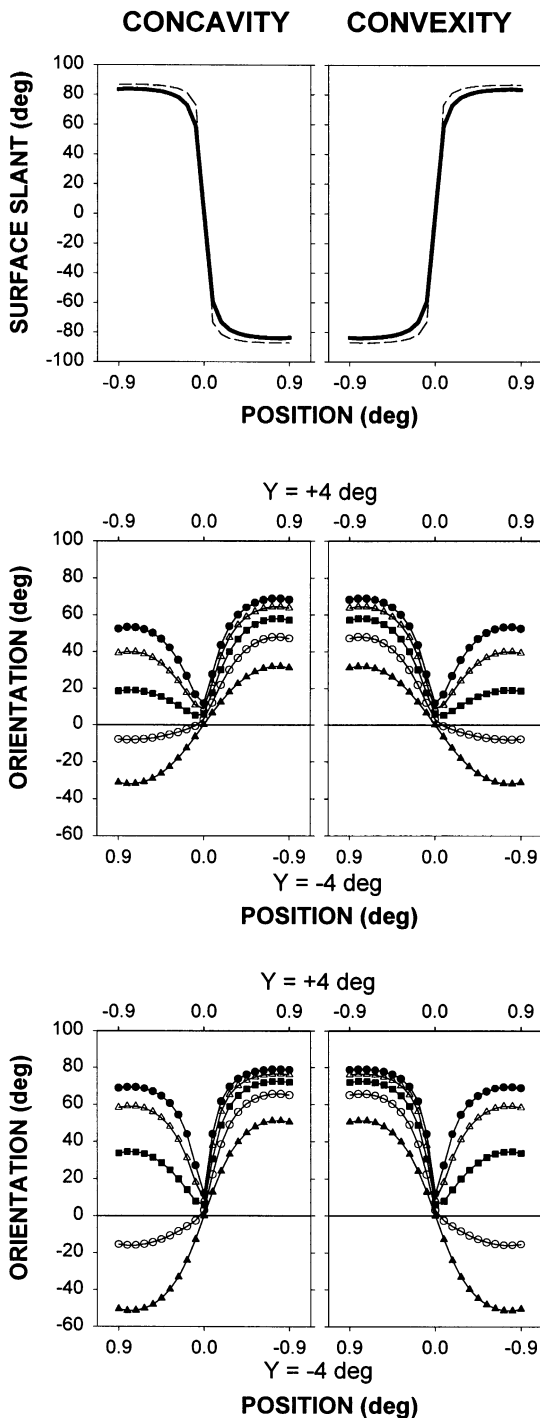


Fig. 8. (a) Data for three observers for the horizontal component. Frequency of reported perceived shape is plotted against simulated shape. The area of each dot represents the frequency with which each perceived shape was reported, with the area totaling unity across each simulated shape category. All observers identified all five shapes correctly. (b) Data for the 22.5° component. Observers could at best distinguish curvatures from slants, but could not distinguish signs of curvatures or directions of slants. (c) Data for the 45° component. The trend is the same as in (b). (d) Data for the 67.5° component. The trend is the same as in (b). (e) Data for the vertical component. The trend is the same as in (b).

was  $1.8^\circ$ . Negative values along the abscissas indicate positions to the left of the mid-point.

In the middle panels of Fig. 9 are plotted the derived orientations for the five components as a function of image position for the sinusoidal surface with an amplitude of one: concavity on the left panel and convexity on the right. The top and bottom abscissas indicate positions across the image at heights of  $+4$  and  $-4^\circ$ ,



respectively. Values for the horizontal component are plotted in solid triangles. If these values were plotted as a linear function of surface slant, they would look like the values for the horizontal component in Fig. 7. The local orientation at the steepest portions of the surface (i.e. slants of  $\pm 84^\circ$  at positions of  $\pm 0.9^\circ$ ) are equal in magnitude but opposite in sign, corresponding to oppositely bowed contours across the top and bottom edges of the images, and thus distinct patterns for projected concavities, convexities, leftward and rightward slants. As the component deviates from horizontal, the asymmetry between orientation values at these slants decreases. For components oriented  $6^\circ$  (solid squares) and more from the axis of maximum curvature, the orientation values at the two slants are of the same sign and steeper than the uncorrugated orientation. This results in positively oblique contours across the top and bottom edges of both the projected concavity and convexity, yielding similar images for the two curvatures. For components oriented  $3^\circ$  (open circles) and less from the axis of maximum curvature, the signs of the orientation values at the two slants take on opposite sign. Thus only texture components oriented within  $3^\circ$  of the axis of maximum curvature will show distinct patterns for the four surface shapes.

The thin dashed lines in the upper panels of Fig. 9, show that doubling the amplitude of the corrugation only increases the magnitude of the surface slant slightly. The derived orientation values for a sinusoidal surface with an amplitude of two are plotted in the bottom panels of Fig. 9. The orientation values at leftward and rightward slanting portions of the surface are steeper than those for the surface with an amplitude of one, however, the component for which the projected orientation values take on opposite signs at the two slants is still approximately  $3^\circ$ . Doubling the amplitude of the corrugation makes little or no difference to this cut-off.

Fig. 9. *Top:* Surface slant as a function of position in the image for a projected concavity (left) and projected convexity (right) of a sinusoidal surface with an amplitude of one (thick, solid lines), and for a sinusoidal surface with an amplitude of two (thin, dashed lines). *Middle:* Projected orientation values for components oriented at  $0^\circ$  (solid triangles),  $3^\circ$  (open circles),  $6^\circ$  (solid squares),  $9^\circ$  (open triangles), and  $12^\circ$  (solid circles) plotted as a function of position in the image, for a sinusoidal surface with an amplitude of one. The top and bottom abscissas indicate image positions across image heights of  $+4$  and  $-4^\circ$ , respectively. For components within  $3^\circ$  (open circles), the local orientation at the leftward and rightward slanting portions of the surface have opposite sign. *Bottom:* Projected orientation values for the same components for a sinusoidal surface with an amplitude of two. The format is identical to the panels in the middle row. Again, for components within  $3^\circ$  (open circles), the local orientation at the leftward and rightward slanting portions of the surface have opposite sign.

5.1. Psychophysical method

Projections of isolated concave, convex, leftward and rightward slanting portions of the surfaces together with flat-frontoparallel surfaces were presented for these five oriented components in the same 5AFC paradigm as that in Experiment 1. A total of 150 conditions (5 component patterns  $\times$  5 shapes  $\times$  3 carrier phases  $\times$  2 depth amplitudes) at ten trials each were randomly interleaved for a total of 1500 trials. These were divided into four separate sessions, each lasting approximately 20 min. Observers AL and FG participated in this experiment.

5.2. Results

Fig. 10 plots for the two observers percent correct identification as a function of the orientation of the component for each of the four shapes: concavity (filled triangles, solid line), convexity (open triangles, dashed lines), rightward slant (open circles, solid line) and leftward slant (filled squares, dashed line). Data for the surface with an amplitude of one are plotted in the left column, for the surface with an amplitude of two in the right column. The dashed horizontal line in each panel indicates chance level performance for a 5AFC task.

For the surface with an amplitude of one, observer AL performed above chance for all four surface shapes only for components oriented within 3° of the axis of maximum curvature. Once the component orientation exceeded 3°, performance for the rightward slant fell below chance. Once it exceeded 6°, performance for the convexity and the leftward slant also fell below chance. Performance for the concavity remained above chance for components oriented out to 12°. Data for the surface with an amplitude of two for this observer also show that performance for all four surface shapes fell above chance only for components oriented within 3° of the axis of maximum curvature. Performance for the concavity fell off for components beyond 3°. Performance for the leftward slant fell off for components beyond 6°, and performance for convexities and rightward slants remained above chance for components out to 12°. Data for observer FG show similar patterns for both surfaces. Performance for all four shapes fell above chance only for the 0° component. For components beyond 0°, performance for the concavity and convexity dropped off. For components oriented beyond 6°, performance fell off for the rightward slant for the surface of amplitude of one. Performance for the leftward slant remained above chance for components out to 12° for this surface. For the surface with an

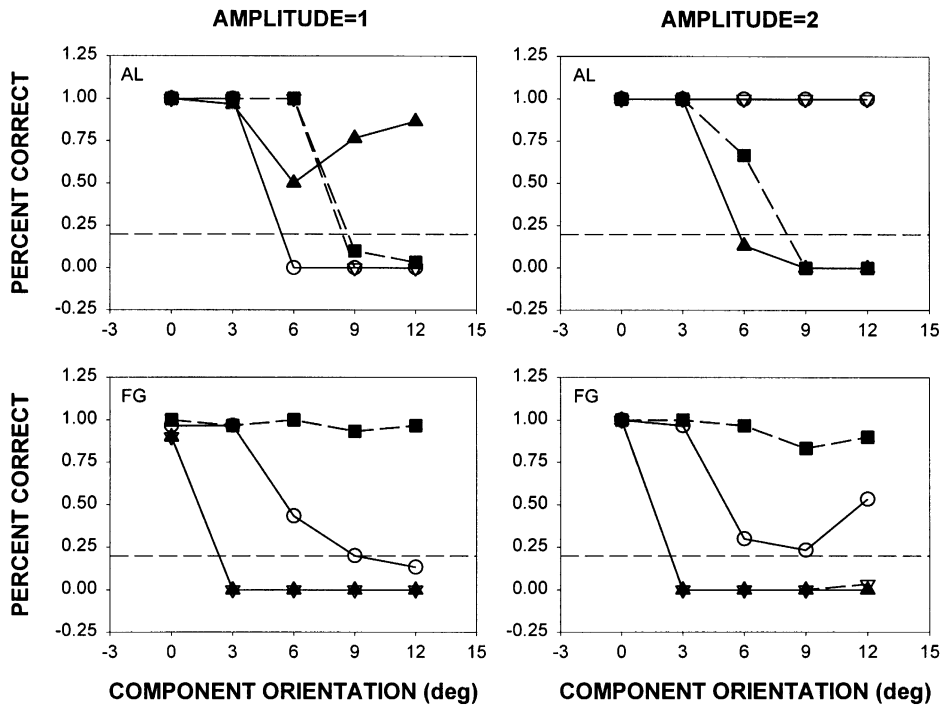


Fig. 10. Percent correct for a 5AFC shape identification task as a function of the orientation of the texture component. Results for each of four surface shapes are plotted in different symbols: concavity (filled triangles, solid line), convexity (open triangles, dashed line), rightward slant (open circles, solid line) and leftward slant (filled squares, dashed line). The dashed horizontal line in each panel indicates chance level performance. Data for the two observers are plotted by row, for a sinusoidal surface with an amplitude of one in the left column, and for a sinusoidal surface with an amplitude of two in the right column. For both surfaces, observer AL perceives all surface shapes correctly only for components oriented within 3° of the axis of maximum curvature. For observer FG, the cut-off falls between 0 and 3°.




Fig. 11. (a) Corrugated horizontal grating added to an uncorrugated vertical grating. The surface appears flattened compared to Fig. 1A, especially along the horizontal mid-line. (b) Corrugated horizontal component of the octotropic plaid added to the remaining seven components uncorrugated. The horizontal component is nearly invisible, and the surface appears nearly flat.

amplitude of two, performance for both rightward and leftward slants remained above chance for components out to  $12^\circ$ . The critical cut-off for this observer falls somewhere between  $0$  and  $3^\circ$  and is unaffected by a doubling of the amplitude of the surface. Note that the two observers have different biases for guessing different shapes; AL tended to perceive convexities and rightward slants correctly while FG tended to perceive leftward and rightward slants correctly.

These results show that for the sinusoidal surfaces used here, only components within about  $3^\circ$  of the axis of maximum curvature will convey the correct shape of the surface. While there are small changes in shape biases, a doubling of the amplitude of the corrugation makes no difference in this critical cut-off.

## 6. Discussion

We have shown formally that for a sinusoidal surface with an amplitude of one, components of a surface texture oriented within  $3^\circ$  of the axis of maximum curvature are the only ones that exhibit distinct 2D patterns of orientation modulations for concavities, convexities, rightward slants and leftward slants. Components oriented along all other lines of curvature exhibit orientation and frequency modulations that are distinct for curvatures on the one hand and slants on the other, but that are nearly identical for both curvatures, and nearly identical for both slants.

The results show that observers utilize the different information provided by texture components along critical and non-critical orientations. For the horizontal and near-horizontal components, which exhibited distinct patterns for concavities, convexities, and leftward and rightward slants, observers correctly identified all surface shapes. For all other components, which exhibited distinct patterns for curvatures and slants, but

similar patterns for both curvatures and for both slants, observers were able to distinguish curvatures from slants, but were unable to distinguish the signs of the curvatures or the directions of the slants.

Stevens (1981) showed that local 3D surface orientation of a developable surface can be computed from projections of parallel rulings along lines of maximum curvature. He suggested that the visual system assumes that surface markings follow projected lines of maximum curvature, and thus rulings along non-maximal lines of curvature result in non-veridical 3D shape percepts. The results show that this assumption is not necessary; the reason that only rulings along projected lines of maximum curvature provide veridical 3D shape percepts is that these are the only ones that provide sufficient information for distinguishing different shape components. Rulings along all other lines of curvature simply do not contain sufficient information.

Although particular patterns of orientation modulations contain all the information required for shape identification, frequency modulations also affect the quality of the shape percept. While the effect is not measured quantitatively in this paper, Figs. 1 and 2 provide a qualitative description. Fig. 1A and B, which contain the requisite patterns of both orientation and frequency modulations, convey surfaces that appear triangular, like folding screens. In contrast, the surfaces in Fig. 1C and D, which contain only the requisite pattern of orientation modulations, appear more curved than triangular, and significantly flattened along the horizontal mid-line of the image where there are no perspective cues. Similar effects hold for the analogous components of the octotropic plaid pattern in Fig. 2.

It is possible that the effect of the vertical components of the texture is simply to line up corresponding points of curvature, e.g. all points along projected lines of maximum curvature falling at a peak of the corrugation. In effect, this would continue the perceived corrugation across the horizontal mid-line of the image which is not affected by perspective projection. However, the images in Fig. 11 show that the effect of the frequency modulations is more substantial. In Fig. 11a, the vertical component of the horizontal-vertical plaid is uncorrugated, i.e. consistent with the projection of a flat surface. While there are no frequency modulations in the image, there are vertical contours that can be used to line up corresponding points of surface curvature. The shape of the surface, however, resembles that conveyed by an image without frequency modulations (Fig. 1C), appearing more curved than the triangular appearance conveyed in Fig. 1A, in which the vertical component is consistent with the corrugation. Similarly in Fig. 11b, the seven non-horizontal components of the octotropic plaid are uncorrugated, consistent with a flat surface, while the horizontal component is corrugated. The curved horizontal component is hardly visi-

ble and the image hardly conveys any 3D shape at all. These contours are nearly invisible because the saliency is similar along all oriented components of the pattern, unlike Fig. 2A in which increases in spatial frequency with increasing slant reduce the saliency along the seven non-horizontal components relative to that of the horizontal component. Frequency modulations thus play a role in the inference of 3D shape, even if they are insufficient for conveying veridical shape in isolation.

Since all the information required for shape identification is contained in the orientation modulations of a narrow band of components, a simple neural model of shape from texture can focus on extracting just these template patterns of orientation modulations. The outputs of local orientation and frequency-selective V1 neurons (Webster & DeValois, 1985; Kulikowski & Vidyasagar, 1986; Parker & Hawken, 1988; Hamilton, Albrecht, & Geisler, 1989; Silverman, Grosz, DeValois, & Elfar, 1989) can be combined to extract patterns of template orientation modulations, such as those in the top row of Fig. 5. Preliminary results show that, using a discrete number of such templates, a matched-filtering algorithm can successfully identify the axis of curvature, the amplitude of the depth corrugation, and the locations of concave and convex curvatures across images of a surface containing several cycles of corrugation (Zaidi & Li, 2000).

In previous models of shape from texture (e.g. Blake et al., 1993; Knill, 1998a,b,c) the observer makes the assumption that flat texture patterns are isotropic (with no bias in orientation) and homogeneous (spectrally similar across translations) (Knill, 1998a,b,c). Where deviations from isotropy and homogeneity occur in the image in the form of gradients such as compression or density, the observer attributes these to the shape of the 3D surface. Although the focus of much of the shape from texture literature has been on quantifying the capacity with which the various texture gradients convey 3D shape and surface slant (Braunstein & Payne, 1969; Gillam, 1970; Vickers, 1971; Rosinski & Levine,

1976; Cutting & Millard, 1984; Todd & Akerstrom, 1987; Blake et al., 1993; Cummings et al., 1993; Goodenough & Gillam, 1997; Knill, 1998a,b,c), the results of Li and Zaidi (2000) together with the results of this paper show that the extraction of texture elements and gradients is neither necessary nor sufficient. The fact that many natural texture patterns cannot be easily segmented into individual texture elements has led to spatial-frequency based approaches to the problem of shape from texture (Bajcsy & Lieberman, 1976; Turner et al., 1991; Sakai & Finkel, 1993; Krumm & Shafer, 1994). However, our results show that modulations in spatial frequency simply do not contain sufficient information to differentiate between curvatures of multiple signs, and are thus insufficient for conveying veridical shape percepts of surfaces containing curvatures of more than one sign. Additionally, since template patterns of orientation modulations can directly provide local surface shape (Zaidi & Li, 2000), this obviates the need to compute affine deformations of the texture pattern and relate them to local surface orientation (Garding, 1992; Malik & Rosenholtz, 1997).

Even though the use of corrugated surfaces was crucial in going beyond the theories developed on the basis of using cylinders or planar surfaces, it will be important in future work to test whether these formal results apply to more complex surfaces. One of the strongest assets of a neural model based on patterns of orientation modulation is that it requires no prior assumptions about the texture pattern itself. This property is likely to be even more useful for non-developable surfaces, such as those made by stretching materials like rubber, where the deformation of the surface alters the texture on the surface. Zaidi & Li (2000) demonstrated that when patterns of orientation modulations like Fig. 1C or D are added to their 90° rotated versions, the results (Fig. 12a,b) convey the shapes of depth plaids, i.e. sums of orthogonal sinusoidal corrugations. However, similar combinations of frequency modulations like Fig. 1E or F and their rotated versions do not convey such shapes (Fig. 12c,d).




Fig. 12. (a) Image of the orientation modulations of the horizontal component on a surface containing a central concavity (Fig. 1C) added to the same image rotated 90°. The resulting image conveys the shape of a depth-plaid, i.e. a sum of orthogonal sinusoidal depth corrugations, containing a central concavity. (b) The same as (a) for a surface containing a central convexity (Fig. 1D added to its rotated version). (c) Image of the frequency modulations of the vertical component on a surface containing a central concavity (Fig. 1E) added to the same image rotated 90°. The resulting image appears flat. (d) The same as (c) for a surface containing a central convexity (Fig. 1F added to its rotated version).

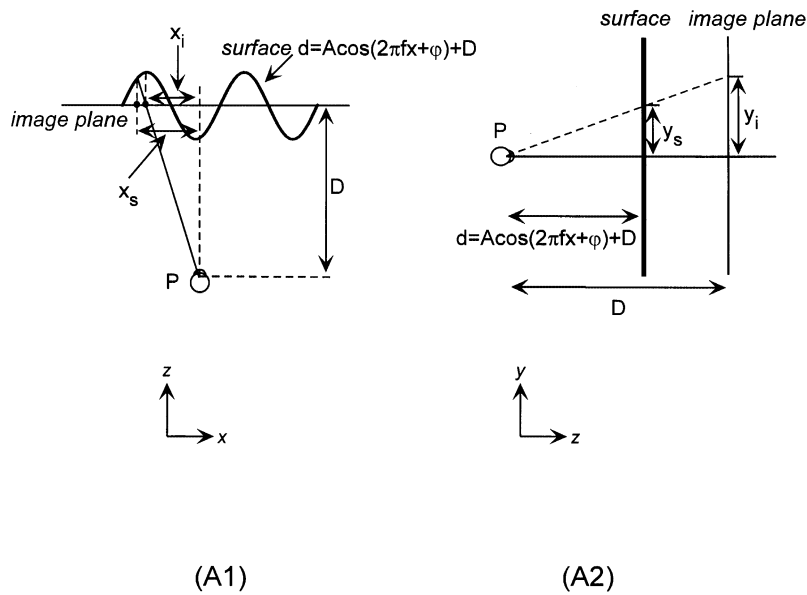


Fig. 13. (A1) Perspective projection in the  $x-z$  plane of a surface corrugated sinusoidally in  $z$  as a function of  $x$ . Each location on the surface is projected onto the image plane by connecting it with  $P$  (the nodal point of the eye). (A2) Perspective projection for a given value of  $x$  of the same surface in the  $y-z$  plane. Since the surface is corrugated in depth ( $z$ ) as a function of  $x$  and not  $y$ ,  $d$  is equal for all values of  $y$ .

**Acknowledgements**

The authors wish to thank A. Fuzz Griffiths and Jeanette Meng for serving as observers, and A. Fuzz Griffiths, Jeanette Meng, Byung-Geun Khang, and Rocco Robilotto for helpful comments on the manuscript. This work was supported by Schnurmacher Institute for Vision Research Award 99-00-87 to Andrea Li, NEI grants EY07556 and EY13312 to Qasim Zaidi, and was presented at the ARVO '00 conference in Ft. Lauderdale.

**Appendix A. Perspective projection of a corrugated surface**

The projection of the corrugated surface is shown in the  $xz$ -plane in Fig. 13. The nodal point of the eye  $P$  is defined as the origin.  $D$  is the distance between  $P$  and the image plane. (Note that this figure is not drawn to scale: in the experiments,  $D$  was 44 cm and the amplitude of depth modulation was 8 cm from peak to trough). In perspective projection, each surface location is projected onto the image plane via rays connecting it with  $P$ .

For each location across the image  $x_i$ , we computed  $x_s$ , the  $x$ -coordinate of the intersection of: (1) the ray connecting  $x_i$  with  $P$  ( $z = Dx_s/x_i$ , where  $P$  lies at  $x = 0$ ,  $z = 0$ ); and (2) the corrugated surface ( $z = A \cos(2\pi fx_s + \varphi) + D$ , where  $A$  is the amplitude,  $f$  is the frequency, and  $\varphi$  is the phase of the depth corrugation).

The intersection is found by equating the two expressions:

$$\frac{Dx_s}{x_i} = A \cos(2\pi fx_s + \varphi) + D \tag{A1}$$

The value of  $x_s$  was then used to compute the length traversed along the surface using the integral:

$$L = \int_0^{x_s} \sqrt{1 + \left(\frac{dz}{dx}\right)^2} dx \tag{A2}$$

where

$$\frac{dz}{dx} = -2\pi fA \sin(2\pi fx + \varphi) \tag{A3}$$

$L$  was used as an index along  $x$  into the (flat) luminance pattern. We then computed the index along  $y$  into the pattern. The depth of the corrugated surface varied only as a function of  $x$ . Therefore, for each value of  $x_i$  the depth of the surface ( $d$ ) was constant for all values of  $y$ . For each value of  $y_i$ , we used  $x_s$  (computed in Eq. (A1)) to compute  $y_s$ , the  $y$ -coordinate of the intersection of the projecting ray connecting  $P$  and  $y_i$ :

$$y_s = \frac{Dy_i}{D + d} \tag{A4}$$

where

$$d = A \cos(2\pi fx_s + \varphi) \tag{A5}$$

The luminance value at the point  $(x_i, y_i)$  in the image was set equal to the luminance at the point  $(L, y_s)$  in the flat texture pattern.

## References

- Bajcsy, R., & Lieberman, L. (1976). Texture gradient as a depth cue. *Computer Graphics and Image Processing*, 5, 52–67.
- Blake, A., Bulthoff, H. H., & Sheinberg, D. (1993). Shape from texture: ideal observers and human psychophysics. *Vision Research*, 33, 1723–1737.
- Braunstein, M. L., & Payne, J. W. (1969). Perspective and form ratio as determinants of relative slant judgments. *Journal of Experimental Psychology*, 81(3), 584–590.
- Cummings, 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*, 113, 196–216.
- Garding, J. (1992). Shape from texture for smooth curved surfaces. *Journal of Mathematical Imaging and Vision*, 2, 327–350.
- Gibson, J. J. (1950). The perception of visual surfaces. *American Journal of Psychology*, 63, 367–384.
- Gillam, B. (1970). Judgments of slant on the basis of foreshortening. *Scandinavian Journal of Psychology*, 11, 31–34.
- Goodenough, B., & Gillam, B. (1997). Gradients as visual primitives. *Journal of Experimental Psychology: Human Perception and Performance*, 23(2), 370–387.
- Hamilton, D. B., Albrecht, D. G., & Geisler, W. S. (1989). Visual cortical receptive fields in monkey and cat: spatial and temporal phase transfer function. *Vision Research*, 29(10), 1285–1308.
- Knill, D. C. (1998a). Surface orientation from texture: ideal observers, generic observers and the information content of texture cues. *Vision Research*, 38(11), 1655–1682.
- Knill, D. C. (1998b). Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Research*, 38(11), 1683–1711.
- Knill, D. C. (1998c). Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture. *Vision Research*, 38, 2635–2656.
- Krumm, J., & Shafer, S. A. (1994). Segmenting textured 3D surfaces using the space/frequency representation. *Spatial Vision*, 8(2), 281–308.
- Kulikowski, J. J., & Vidyasagar, T. R. (1986). Space and spatial frequency: analysis and representation in the macaque striate cortex. *Experimental Brain Research*, 64(1), 5–18.
- Li, A., & Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research*, 40, 217–242.
- 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.
- Marroquin, J. L. (1976). *Human visual perception of structure*. Unpublished M.S., M.I.T., Cambridge.
- Mundy, J. L., & Zisserman, A. (1992). Projective geometry for machine vision. In J. L. Mundy, & A. Zisserman, *Geometric invariance in computer vision* (pp. 463–519). Cambridge, MA: M.I.T. Press.
- Parker, A. J., & Hawken, M. J. (1988). Two-dimensional spatial structure of receptive fields in monkey striate cortex. *Journal of Optical Society of America A*, 5(4), 598–605.
- Reichel, F. D., & Todd, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology*, 16(3), 653–664.
- Rosinski, R. R., & Levine, N. P. (1976). Texture gradient effectiveness in the perception of surface slant. *Journal of Experimental Child Psychology*, 22, 261–271.
- Sakai, K., & Finkel, L. H. (1993). Characterization of the spatial-frequency spectrum in the perception of shape from texture. *Journal of Optical Society of America A*, 12(6), 1208–1224.
- Silverman, M. S., Grosz, D. H., DeValois, R. L., & Elfar, S. D. (1989). Spatial-frequency organization in primate striate cortex. *Proceedings of the National Academy of Science USA*, 86(2), 711–715.
- Stevens, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence*, 17, 47–73.
- Todd, J. T., & Akerstrom, R. A. (1987). Perception of three-dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance*, 113, 221–224.
- Todd, J. T., & Reichel, F. D. (1990). Visual perception of smoothly curved surfaces from double-projected contour patterns. *Journal of Experimental Psychology*, 16(3), 665–674.
- Turner, M. R., Gerstein, G. L., & Bajcsy, R. (1991). Underestimation of visual texture slant by human observers: a model. *Biological Cybernetics*, 54, 215–226.
- Vickers, D. (1971). Perceptual economy and the impression of visual depth. *Perception and Psychophysics*, 10(1), 23–27.
- Webster, M. A., & DeValois, R. L. (1985). Relationship between spatial-frequency and orientation tuning of striate-cortex cells. *Journal of Optical Society of America A*, 2(7), 1124–1132.
- Zaidi, Q., & Li, A. (2000). Neural model of shape from texture: developable surfaces. *Investigative Ophthalmology and Visual Science (Supplement)*, 41, S219.