



PERGAMON

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Vision Research 43 (2003) 2741–2758

Vision
Research

www.elsevier.com/locate/visres

Observer strategies in perception of 3-D shape from isotropic textures: developable surfaces

Andrea Li ^{*}, Qasim Zaidi

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

Received 4 February 2003; received in revised form 25 June 2003

Abstract

We document the limitations of isotropic textures in conveying three-dimensional shape. We measured the perceived shape and pitch of upright and pitched corrugated surfaces overlaid with different classes of isotropic textures: patterns containing isotropic texture elements, isotropically filtered noise patterns, and patterns containing ellipses or lines of all orientations. Frequency modulations arising from surface slant were incorrectly interpreted as changes in surface distance, resulting in concavities being misclassified as convexities, and right and left slants as concavities. In addition, images of pitched surfaces exhibited oriented flows that confound surface shape and surface pitch. Observers related oriented flow patterns to particular surface shapes with a bias for perceiving convex surfaces. When concave and convex curvatures were concurrently visible, the number of correct shape classifications increased slightly. Isotropic textures thus convey correct 3-D shapes of developable surfaces only in some conditions, and the same perceptual strategies lead to non-veridical percepts in other conditions.

© 2003 Elsevier Ltd. All rights reserved.

Keywords: Shape from texture; Isotropic texture; Frequency modulations; Orientation modulations

1. Introduction

Isotropic textures are those whose amplitude spectra contain equal energy at all orientations. Some examples are shown in Fig. 1. These textures can consist of isotropic elements such as the patterns in Fig. 1A, they can consist of isotropically filtered noise that cannot be easily segmented into individual elements such as the patterns in Fig. 1B, or they can consist of oriented elements but where the set of elements span all orientations such as the patterns shown in Fig. 1C and D. When an isotropic texture is overlaid onto a 3-D surface that is then projected to a perspective image, the texture in the image deviates from isotropy, and these deviations contain information about the 3-D shape. As such, numerous psychophysical studies have utilized isotropic textures in the study of shape from texture (e.g. Cummings, Johnston, & Parker, 1993; Cutting & Millard, 1984; Knill, 1998a).

The surfaces used in these studies were either flat and slanted out of the fronto-parallel plane or singly curved

like a cylinder. Both kinds of surfaces are developable, i.e. they are formed by bending a piece of paper and leaving the texture on the surface unaltered. For the present study, we used a developable surface that was sinusoidally corrugated in depth containing both concave and convex portions.

Li and Zaidi (2000, 2001b) showed that the use of an upright corrugated surface critically separated surface textures into two groups. The first group consists of textures that contain specific patterns of orientation modulations in perspective images of concave, convex, rightward and leftward slanting surface shapes. In psychophysical experiments, observers correctly classified surface shapes overlaid with these textures. Images of surfaces overlaid by the second group of textures do not contain the requisite orientation modulations, but do contain frequency modulations along the axis of maximum curvature. The second group includes textures that have isotropic frequency spectra. Li and Zaidi (2000) and Zaidi and Li (2002) demonstrated this using isotropically filtered noise patterns and polka-dot patterns overlaid onto half-cycles of a corrugated surface. Vertical spatial frequencies in the image were lowest at the center of the projected concavities and convexities where

^{*} Corresponding author. Fax: +1-212-780-5137.

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

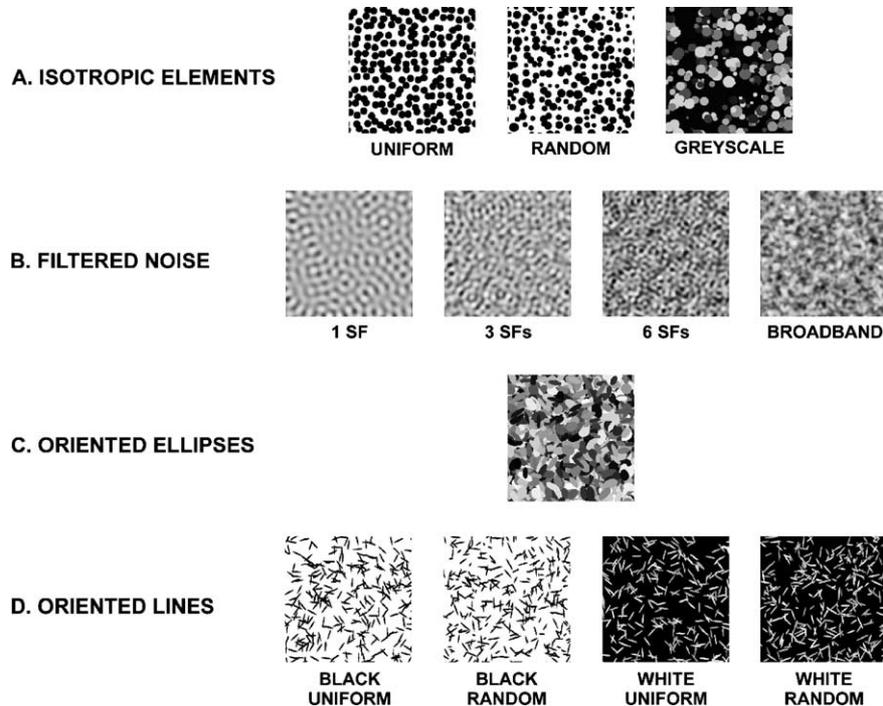


Fig. 1. Isotropic patterns used in our experiments. (A) Patterns consisting of individual isotropic elements, uniform in size, random in size, or random grey-levels and overlapping. (B) Isotropically filtered noise patterns containing one, three or six spatial frequencies, or a broadband range of frequencies. (C) Pattern of grey-level oriented ellipses the orientations of which are chosen from an isotropic range. (D) Patterns of oriented lines, uniform or random in length, either black on a white background or white on a black background. The orientations are chosen from an isotropic range.

the surface was fronto-parallel with respect to the observer, and highest along the leftward and rightward slanting portions of the surface. Concave half-cycles thus resembled convex half-cycles, and rightward slanting half-cycles resembled leftward slanting half-cycles. Rather than interpret these differences in spatial frequency as differences in surface slant, observers interpreted them as differences in distance from the surface with low frequencies interpreted as closer to the observer rather than as fronto-parallel portions of the surface, and high frequencies as farther rather than as slanted portions. As a result, concavities and convexities were both classified as convex, and rightward and leftward slanting portions were classified as concave.

In the present study, we extend the results of Li and Zaidi (2000) and Zaidi and Li (2002) to different classes of isotropic textures and different surface orientations. Our results will confirm that for upright developable surfaces overlaid with various isotropic textures, observers perceive 3-D shapes as if frequency modulations in perspective images were due to changes in distance rather than slant. The results will show that this strategy is also used for developable surfaces pitched out of the fronto-parallel plane, but that an additional strategy is to identify particular patterns of oriented flows with particular surface shapes at particular pitches. When the same surface is viewed such that multiple surface curvatures are visible, the number of correct classifications

is slightly increased suggesting that surrounding surface curvatures can help in the perception of 3-D shape from isotropic textures. However, even when multiple curvatures are visible, isotropic textures convey non-veridical percepts about as often as they convey veridical ones.

2. Spatial frequency in the perspective image

To illustrate how frequency modulations arise in perspective images of 3-D surfaces, we will start with a fronto-parallel planar surface. The spatial frequency in the perspective image of this flat surface depends not only on the distance between the surface and the observer, but also on the orientation of the surface in 3-D space (see Appendix of Zaidi & Li (2002)). For example, as shown in the top row of Fig. 2, the frequency in the perspective image of a surface overlaid with a vertical sinusoidal grating will be constant across the image. If the surface is moved farther away from the observer, the frequency in the image will increase uniformly. The second row in Fig. 2 shows the same surface viewed from 1.1, 1.3, and 2.1 times the distance respectively from left to right. If instead the surface is slanted to the left by 20°, 40° and 60° (third row, Fig. 2) so that the right edge of the surface is closer to the observer than the left edge, or to the right by the same amounts (bottom row, Fig. 2), the perspective image exhibits a

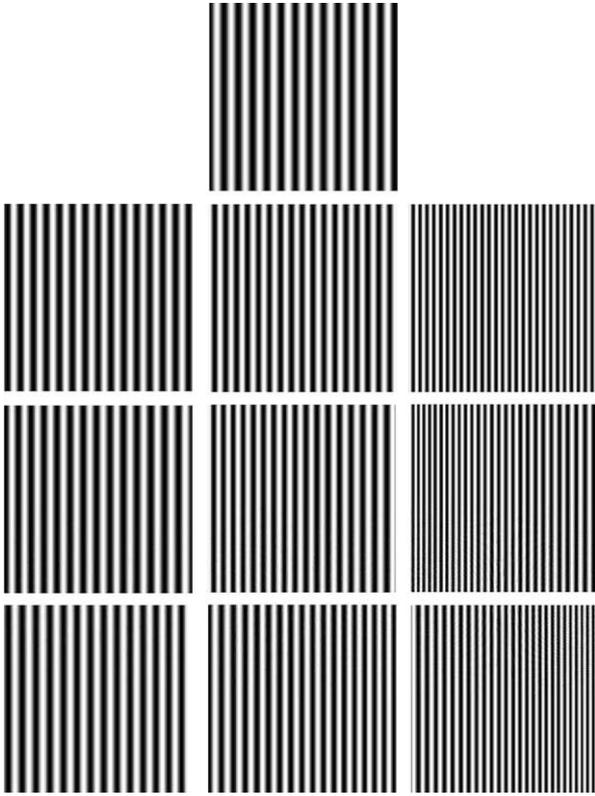


Fig. 2. Spatial frequency in the projected image varies as a function of distance and slant. The panel in the top row shows a flat fronto-parallel surface overlaid with a vertical grating, computed to span 9° viewed at 1 m. Frequency increases uniformly across the image as the viewing distance is increased to 108, 131, and 208 cm as shown in the second row. Frequency also increases if the surface is viewed at 1 m, but slanted to the left by 20° , 40° , and 60° (third row), or to the right by 20° , 40° , and 60° (bottom row).

change in the local frequency across the image with low frequencies along the closer edge of the surface and high frequencies along the farther edge. The detectability of the frequency modulations depends on both the base frequency and the amount of modulation brought about by the surface slant (see Jamar, Campagne, & Koenderink, 1982). For moderate slants (e.g. middle panels of third and fourth rows in Fig. 2) and for small windows of an extremely slanted surface, the frequency modulation will be less detectable and instead the frequency will simply appear to be increased uniformly across the surface, regardless of whether the surface is slanted to the left or to the right. Thus increasing the distance of the surface and increasing the magnitude of the slant angle have similar effects on the frequency information in the image, and the visual system must determine the source of these changes when determining the 3-D shape of the surface. For the parameters used in our experiments, the changes in depth along the surface are small relative to the viewing distance. As a result, frequency modulations in the image are largely a result of surface slant (see Appendix A).

Fig. 3 shows surfaces overlaid with an isotropic polka-dot pattern. An increase in distance reduces the sizes of the elements (isotropic frequency increases in second row), but in the image, slants cause frequency transformations along the horizontal axis and distort the polka dots in the image. Specifically, the dots are isotropic in the image where the surface is fronto-parallel, and vertically elliptical where the surface is slanted to the left or to the right (third and fourth rows of Fig. 3). Slight changes in the vertical extent will occur from changes in distance. However, for the surfaces used in our experiments the effect of distance on frequency will be small. At moderate slants and/or for small windows, the dots do not contain sufficient information to correctly distinguish these two slants and thus to correctly convey the shape of a vertically corrugated surface. Changes in surface slant will thus result in a change in shape of the projected dots along the axis of slant or curvature in the image, whereas changes in viewing distance will result in an isotropic change in size. How observers interpret changes in spatial frequency in the projected image in perceiving 3-D shape is the first empirical question we address in this paper.

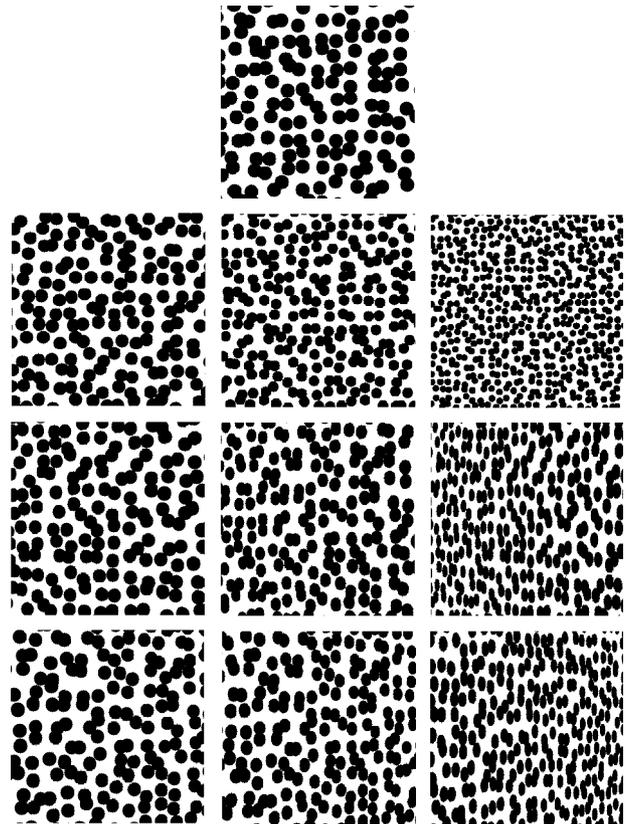


Fig. 3. The same flat surface at the same viewing distances and slants as in Fig. 2, but overlaid with an isotropic polka-dot pattern. Notice that as distance increases (second row), the sizes of the dots decrease isotropically, while as slant increases to the left or right (third and bottom rows), it is predominantly the horizontal spans of the dots that decrease (along with some decrease in vertical span due to distance).

For a vertically corrugated surface pitched out of the upright orientation (towards a ground or ceiling plane), the components of a texture aligned with other axes of surface curvature will exhibit changes not only in frequency but also in orientation that depend on the shape and distance of the surface, and also on its pitch (Zaidi & Li, 2002). Thus for developable surfaces in orientations other than upright, the available information in the perspective image of the surface textured with an isotropic texture will be more complicated than simple frequency modulations along a single axis. How observers use the combined information in images of pitched surfaces to make judgments of perceived shape and perceived pitch is the second empirical question we address in this paper.

The purpose of Experiment 1 is to examine the capacity with which different isotropic textures convey the 3-D shapes of developable surfaces at various pitches. Experiment 2 examines the perceived pitch of the same surfaces used in Experiment 1. In Experiment 3, we examine the effect on perceived shape of enlarging the view of the same surfaces such that both concave and convex curvatures are visible.

3. Experiment 1: perceived shape

In this experiment, observers were asked to judge the 3-D shape of four different half-cycles of a corrugated surface that were overlaid with various isotropic textures and projected in perspective.

3.1. Stimuli

Four classes of isotropic texture patterns were used in the experiment (Fig. 1). The first class (Fig. 1A) consisted of patterns made up of discrete isotropic dots. The dots were either equal in size ($15'$ radius), random in size within a specified range ($8'$ – $20'$ radius), or random in size ($4'$ – $25'$ radius) of different grey-levels and overlapping. The second class (Fig. 1B) consisted of isotropic noise patterns generated by filtering white noise with isotropic filters. The patterns contained a single spatial frequency (1.5 cpd), three frequencies (0.75, 1.5, 2 cpd), six frequencies (0.5, 0.75, 1, 1.5, 2, 2.5 cpd), or a broadband range of frequencies (Gaussian low-pass with standard deviation of 1.46 cpd). The third class (Fig. 1C) consisted of ellipse-shaped elements whose orientations were randomly chosen from 0° to 360° resulting in a globally isotropic pattern. The sizes of the major and minor axes of each ellipse were chosen randomly ($7'$ – $21'$ for the minor axis, $21'$ – $35'$ for the major axis). The ellipses were different grey-levels and they overlapped so that the background was completely covered. The fourth class consisted of oriented line

elements, black lines on a white background or white lines on a black background, the orientations of which were chosen randomly from 0° to 360° . The lines were either uniform in length ($21'$), or random in length ($7'$ – $28'$). It is worth reiterating that we have used the term isotropic to refer to texture patterns that contain equal spatial energy at all orientations, i.e. whose global amplitude spectra are isotropic, even if the individual elements are not.

Each texture pattern was overlaid onto five different surface shapes shown in Fig. 4. One was flat and the other four were half-cycles of a surface corrugated sinusoidally in depth as a function of horizontal position: a concavity, a convexity, a right slant spanning a near peak to far peak from left to right, and a left slant spanning a far peak to a near peak from left to right. Note that left and right edges of the two slanted surfaces terminated at the centers of a concavity and convexity of the corrugation and as such, the surfaces were fronto-parallel with respect to the observer at those points. Since frequency modulations in the image are caused by slant, both concave and convex half-cycles exhibit similar high–low–high frequency gradients, and both right and left slanted half-cycles exhibit low–high–low gradients. These slant-based frequency similarities would not occur in images of developable surfaces with sharp edges. For example, for a surface resembling an upright folding screen, a sharp “concavity” (i.e. a hinge pointing away from the observer) would not exhibit local low frequencies because of the lack of fronto-parallel portions.

Fig. 5 shows the relative dimensions of the stimulus arrangement. Each half-cycle shape was computed such that the zero-crossings coincided in depth with the image plane, i.e. the peaks of the convexity and concavity were respectively in front of and behind the image plane, and the centers of the left and right slants were at the image plane. The flat surface was coincident with the image plane. The perspective images were computed such that when viewed at 1 m, the retinal images coincided with those of a real 3-D corrugated surface with amplitude of 7 cm and wavelength of 15.4 cm. At this viewing distance, the projected image spanned $9^\circ \times 4.5^\circ$ (15.4×7.7 cm).

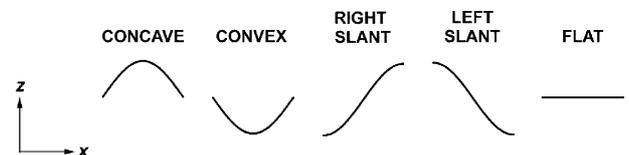


Fig. 4. Five surface shapes used in our experiments. Four half-cycles of a sinusoidal corrugation: a concavity, a convexity, a right slant spanning the fronto-parallel centers of a convexity and concavity, a left slant spanning the fronto-parallel centers of a concavity and convexity; and a flat fronto-parallel surface.

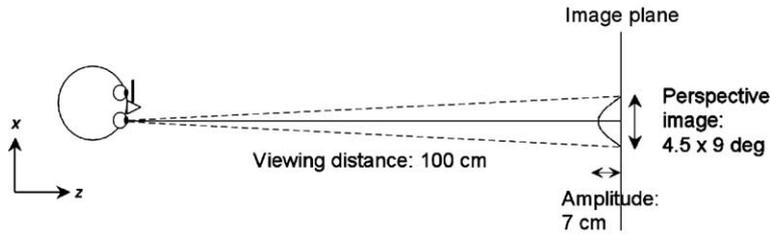


Fig. 5. Dimensions of stimuli and viewing conditions.

Each of the five shapes was presented at one of three surface pitches: 0° (upright), $+30^\circ$ (pitched forward towards the ground plane so that the top edge of the surface was farther from the observer than the bottom edge), and -30° (pitched backward towards the ceiling plane so that the top edge of the surface was closer than the bottom edge).

3.2. Methods

Images were presented on a SONY GDM-F500 flat screen monitor with a 800×600 pixel screen running at a refresh rate of 100 frames/s via a Cambridge Research Systems Visual Stimulus Generator (CRS VSG 2/3) controlled through a 400 MHz Pentium II PC. Through the use of 12-bit DACs, after gamma correction, the VSG was able to generate 2861 linear levels per gun.

Images were interleaved randomly in a 5AFC paradigm in which observers were asked to indicate whether the shape of the surface appeared concave, convex, slanted to the right, slanted to the left, or flat. After the response was made, the next image was automatically presented without delay. The experiment was divided into four sessions. In a single session, observers were presented with images of all five shapes at all three pitches for three different surface patterns. Each image was presented 10 times, resulting in 450 trials per session. Observers initially adapted to a mid-grey screen of 25 cd/m^2 containing a central fixation for 1 min. In each trial, a perspective image of a textured surface containing the central fixation was presented for 1 s accompanied by a short audible beep. After the image disappeared, observers reported the perceived shape of the surface using a three-toggle response box. No feedback was given. Each session lasted approximately 15 min. Viewing was monocular in a dark room with the head stabilized by a chinrest so that the observer's eyes were at the same level as the center of the screen.

Before the experiment, observers were instructed that the surface might sometimes appear pitched forward or backward out of the upright position, but that they should disregard the pitch and simply report the perceived shape of the surface. They were also instructed that sometimes the shape of the surface might not appear symmetrically concave or convex, but that they

should judge asymmetrical concavities and convexities respectively as they would judge symmetrical concavities and convexities, i.e. they should simply judge the relative positions in depth of the left and right edges and the central portion of the image. An initial practice session familiarized each observer with the response box.

Seven naive observers were paid to participate in this study: three students from SUNY College of Optometry, one undergraduate student, and three members of the SUNY staff. All had normal or corrected-to-normal acuity.

3.3. Results

3.3.1. Upright surfaces

Fig. 6 shows the perspective images of the five different upright surface shapes for each of the 12 patterns. For all surface patterns, the projected images exhibit changes in spatial frequency along the axis of maximum surface curvature, i.e. the horizontal axis. Images of the concavity and the convexity both contain lower spatial frequencies along the vertical mid-line and higher frequencies along the left and right sides of the image. Conversely the right and left slants contain higher frequencies along the central mid-line and lower frequencies along the left and right sides because the left and right edges of the slants were fronto-parallel with respect to the observer.

For the oriented line and ellipse patterns, the orientations of the elements modulate across the image, tending towards vertical where the surface is slanted to the right or to the left. This is because the steepness of all oriented components increases as a function of slant (Li & Zaidi, 2001a).

Data averaged across the seven observers for each pattern are presented in Fig. 7. Each plot shows the frequency of perceived shapes for each of the five simulated shapes. The frequency is coded as the area of each dot, with the areas across the five perceived shapes adding up to unity for each simulated shape. If all five shapes were classified correctly on 100% of the trials, the plot would exhibit large dots only along the positive diagonal.

For all 12 patterns, the data show similar trends: concavities and convexities are reported as convex, and

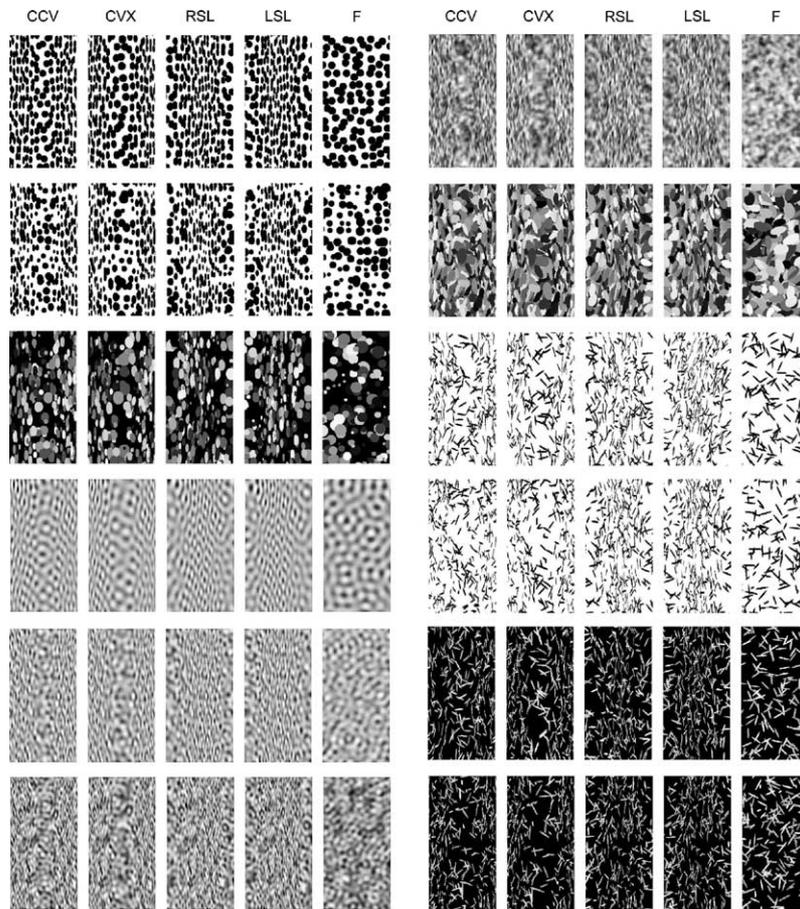


Fig. 6. All surface shapes overlaid with all texture patterns in the upright (pitch = 0) orientation. Concavities and convexities both show decreased frequencies along the central vertical mid-line and both right and left slants show decreased frequencies along the left and right edges of the images.

both slants are reported as concave. These results are consistent with the strategy that frequency in the perspective image is interpreted as distance rather than slant, with lower frequencies being identified with closer (rather than fronto-parallel) portions of the surface and higher frequencies as farther (rather than slanted). For the oriented line patterns, a greater number of surfaces were reported as flat.

3.3.2. Surfaces pitched out of the upright position

Figs. 8 and 9 show images of the same surfaces when they are pitched backward and forward respectively. These figures show that when the surfaces are pitched out of the upright position, in addition to frequency modulations, the isotropic patterns exhibit orientation flows of the kind described by Hel Or and Zucker (1989) and Knill (2001). The frequency modulations are similar to those for upright surfaces: lowest frequencies at the peaks and troughs of the corrugation and highest frequencies along the slanted portions of the surface. The orientation flows, on the other hand, follow four specific patterns that depend not only on the shape of the surface, but also on the pitch. Fig. 8 shows that when the surface is pitched backward so that the top edge is closer

to the observer than the bottom edge, the concavity exhibits flows that are downwardly bowed across the image, while the convexity exhibits upwardly bowed flows. The right slant exhibits positively oblique flows from bottom left to top right and the left slant exhibits negatively oblique flows from top left to bottom right. Fig. 9 shows that when the surface is pitched forward so that the top edge is farther than the bottom edge, these flow patterns reverse within each shape pair, i.e. the concavity exhibits upwardly bowed flows and the convexity exhibits downwardly bowed flows. Similarly, the right and left slant flows are reversed.

Fig. 10 shows perspective images of the four half-cycle surfaces at the same two pitches (+ = forward, - = backward), overlaid with a horizontal grating pattern. Because the surfaces are curved in depth as a function of the horizontal axis, the horizontal contours of the grating are aligned with projected lines of maximum surface curvature, i.e. the depth of the surface changes maximally along these contours. The flow patterns in Figs. 8 and 9 follow these lines of curvature and exhibit the same reversals: concavities and convexities pitched forward respectively exhibit upwardly and downwardly bowed flows that reverse when the surfaces

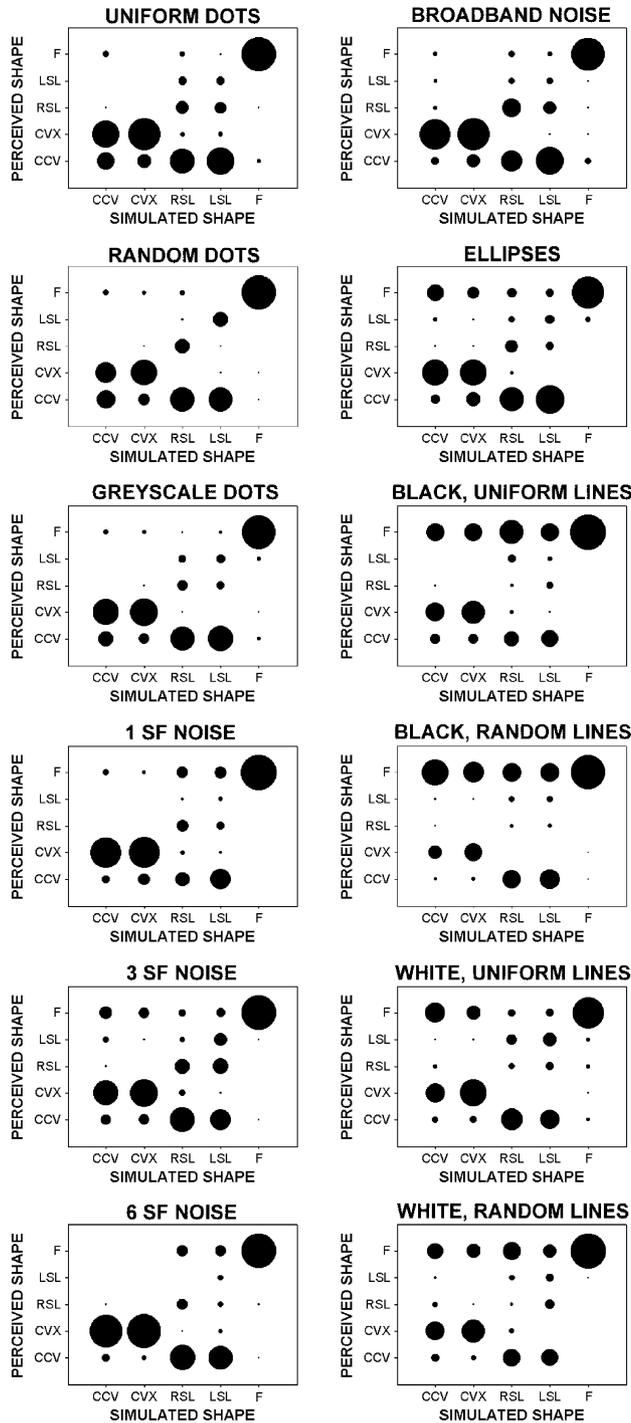


Fig. 7. Averaged data for upright surfaces shown in Fig. 6. Each panel plots for a single pattern the frequency with which each of the five simulated shapes was reported as each of the five perceived shapes. Frequency is coded as the area of each dot. Correct classifications for all five shapes would result in large dots only along the positive diagonal. Observers tend to classify both concavities and convexities as convex and both slants as concave, or they classify all surfaces as flat.

are pitched backward, and similarly the right and left slants pitched forward respectively exhibit negatively and positively oblique flows that reverse when the surfaces

are pitched back. The flows for slanted surfaces and for extrema of curvatures are qualitatively different from one another, but they confound sign of curvature or slant with sign of pitch.

Data for surfaces pitched backward are shown in Fig. 11. For some texture patterns, a substantial proportion of the judgments are correct, as evidenced by the dots along the diagonal (see all polka-dot patterns, broadband noise, ellipses, and black lines). For these patterns, observers appear to be correctly relating the patterns of oriented flows with their respective shapes (see Fig. 10, bottom row). For other texture patterns, a large proportion of judgments reflect the strategy in which frequency is interpreted as distance, so that concavities and convexities are both reported as convex and right and left slants are reported as concave (see noise patterns and white line patterns). For all 12 patterns, the convex image containing upwardly bowed flows was most often classified correctly as convex, and for eight out of the 12 patterns the concave image containing downwardly bowed flows was classified as convex at least as often as concave. As in the upright case, some subjects had difficulty perceiving any 3-D shape for the oriented line patterns and reported all shapes as flat.

Data for surfaces pitched forward are shown in Fig. 12. Overall the data show a much smaller number of correct judgments than those for surfaces pitched backward, especially for the slanted surfaces. Notice that for some patterns (dots, ellipses, black lines) right slants are often misclassified as left slants, and left slants misclassified as right slants. Notice also that for these patterns, concavities and convexities are similarly confused for one another, whereas in the backward pitch condition, convexities were rarely reported as concavities. These reversals suggest that observers are relating particular flow patterns with particular surface shapes. Since the flow patterns are reversed from what they were for surfaces when they were pitched backward, observers are now incorrectly reversing their judgments when the surface is pitched forward. As in the upright and pitched backward cases, some observers had a difficult time perceiving any 3-D shape for the oriented line patterns and reported all shapes to be flat. In addition to these response reversals, a large proportion of the judgments are consistent with the frequency strategy, most notably for concavities and convexities, and most notably for the filtered noise patterns. For 10 out of the 12 patterns, both concavities *and* convexities are most often classified as convex.

The stimuli used in this experiment and the contours in Fig. 10 all exhibit frequency modulations in addition to orientation modulations, so it is difficult to separate frequency effects from orientation effects. Fig. 13 shows the orientation modulations from Fig. 10 without the frequency modulations: the contours within each panel are all identical and equally spaced. (The panel with the

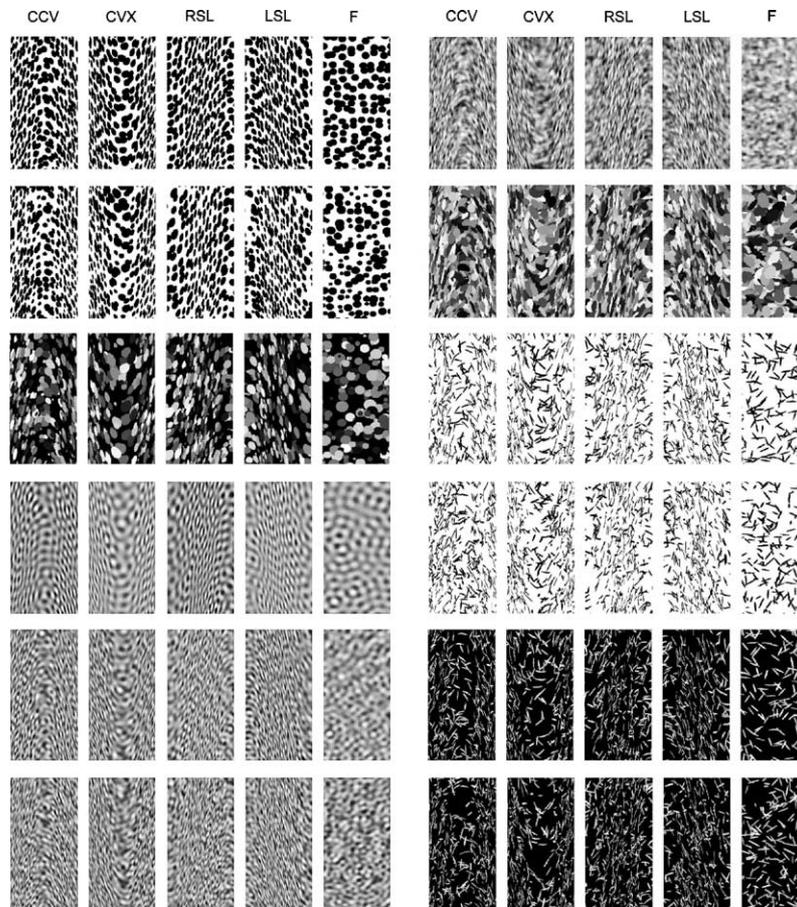


Fig. 8. All surface shapes overlaid with all texture patterns for surfaces pitched backward (pitch = -30°). Images show similar frequency modulations as those for upright surfaces in Fig. 6, but additionally they show patterns of oriented flows that depend on the shape of the surface.

upwardly bowed contours was made by tracing a single upwardly bowed contour from Fig. 10 and duplicating it at equal distances along the vertical axis. The entire panel was then rotated 180° to yield the panel with the downwardly bowed contours. The right two panels were similarly generated by duplicating an oblique contour.) There are no longer any perspective changes, and there are no frequency modulations because the contours are now exactly parallel. Under casual observation, the downwardly and upwardly bowed contours elicit percepts of convex surfaces pitched forward and backward respectively, though the percepts can be bi-stable for some observers. However the oblique contours do not elicit strong percepts of concavities. This suggests that the tendency towards perceiving concavities and convexities as convex likely results from both the bias of perceiving orientation modulations bowed both upwardly and downwardly as convex, and the misinterpretation of frequency modulations as distance. For images containing oblique flows, the concave percepts probably result from the latter. The bias accounts for the fact that, for surfaces pitched forward and backward, observers classified concavities and convexities as convex more often than they classified slants as concave.

Taken together, the data are consistent with two strategies when judging the shape of pitched surfaces. The first strategy is to use the frequency modulations as cues to distance, even though the frequency modulations are along one dimension and thus physically consistent only with changes in surface slant. This strategy biases the perception of both concavities and convexities towards convex, and both right and left slants towards concave. The second strategy is to identify each of the four oriented flow patterns with a specific shape, consistent with a backward pitch of the surface. Since the flow patterns are reversed within shape pairs when the surface is pitched forward, observers' shape judgments are also reversed.

4. Experiment 2: perceived pitch

In Experiment 1, observers were asked to judge the perceived shape of different surfaces regardless of whether or not they appeared pitched out of the fronto-parallel plane. In Experiment 2, we measured the perceived pitch of the same stimuli as those used in Experiment 1. Observers were asked to judge whether

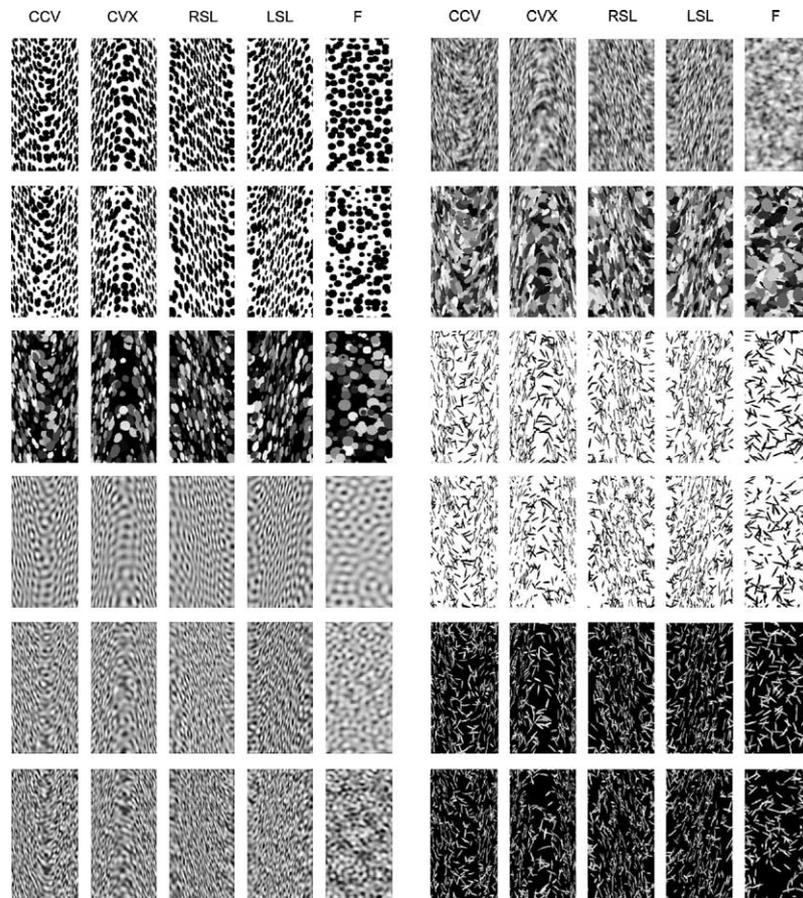


Fig. 9. All surface shapes overlaid with all texture patterns for surfaces pitched forward (pitch = +30°). Images show similar frequency modulations as those for upright surfaces in Fig. 6, but additionally they show patterns of oriented flows that depend on the shape of the surface. Patterns are similar to those in images of surfaces pitched backward in Fig. 7, but reversed within shape pairs, e.g. flows for a concavity pitched forward resemble flows for a convexity pitched backward and vice versa.

the surface appeared upright, pitched forward, or pitched backward, regardless of the surface shape.

4.1. Stimuli

The stimulus set was identical to that used in Experiment 1 so the stimuli can be seen in Figs. 6, 8, and 9.

4.2. Methods

Observers were instructed that the surface might appear upright, pitched forward so that the top edge was farther than the bottom edge, or pitched backward so that the top edge was closer than the bottom edge, and that they should choose the pitch that most closely resembled their percept using the response box. They were also told that the surfaces would appear to take on various 3-D shapes as they had seen in the previous experiment, but that they should try to ignore the shape and simply judge the pitch. Conditions were divided into four sessions and stimulus presentation and viewing

conditions were identical to those in Experiment 1. Each session lasted approximately 15 min.

Four out of the seven observers used in Experiment 1 participated in this experiment.

4.3. Results

Fig. 14 shows the results from this experiment. Each of the 12 rows represents data averaged across all four observers for a single texture pattern. The leftmost panel in each row plots data for the simulated concavity, the middle panel for the simulated convexity, and the rightmost panel for the simulated flat surface. Across all 12 patterns, data for the right and left slants showed no systematic trends and so will not be shown. Each panel represents the frequency with which each of the three simulated pitches (0 = upright, + = forward, and - = back) was reported as each of the three perceived pitches. As in previous data plots, the frequency is represented by the area of each black dot with the area summing to 100% across the three perceived pitches for a single simulated pitch.

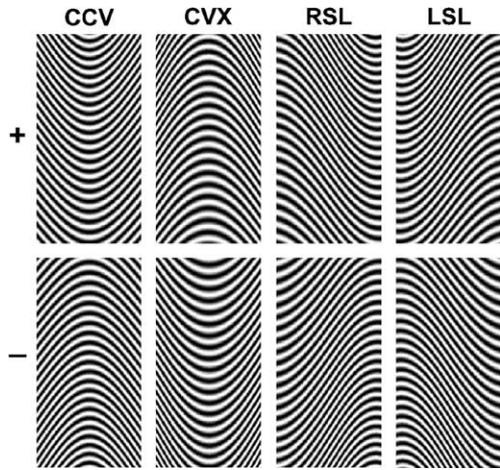


Fig. 10. Projected lines of maximum surface curvature for the same four half-cycle shapes as those used in Experiment 1 pitched forward (+) and backward (-) by 30°. Oriented flows shown in Figs. 8 and 9 follow these lines.

The most obvious trend in the data is for flat surfaces (right panels). Across all patterns, the data show that for the viewing conditions and pitches used in this experiment, flat surfaces appeared upright regardless of the simulated pitch. This suggests that for the present stimulus parameters, isotropic textures are unable to correctly convey surface pitch out of the upright position, despite small changes in frequency along the vertical axis of the image. At greater pitches and/or smaller viewing distances, greater perspective and frequency information could possibly increase the number of correct pitch judgments, however the present data are consistent with previous studies suggesting that observers tend to perceive pitched surfaces as upright (Gibson, 1950; Turner, Gerstein, & Bajcsy, 1991).

Data for the simulated concavity and convexity in the left and middle panels are less consistent. For the 10 patterns other than the white line patterns in the last two rows, data for the simulated positive and negative pitches appear to show a particular trend: for the simulated convexity (middle panels), the forward and backward pitch appear to be classified correctly for many trials, however for the simulated concavity (left panels), the surface pitched forward is classified in many cases as pitched backward, and the surface pitched backward is classified as pitched forward. For both concavities and convexities in the upright position, observers' responses show little consistency.

Fig. 10 shows that a concavity pitched forward (+) and a convexity pitched backward (-) both exhibit upwardly bowed flows. Fig. 14 shows that observers classified both of these surfaces as pitched backward, i.e. the pitch of the convexity. Similarly, a concavity pitched backward and a convexity pitched forward exhibit downwardly bowed flows, and observers classified both of these as pitched forward, i.e. again the pitch of the convexity. These identifications were made despite the

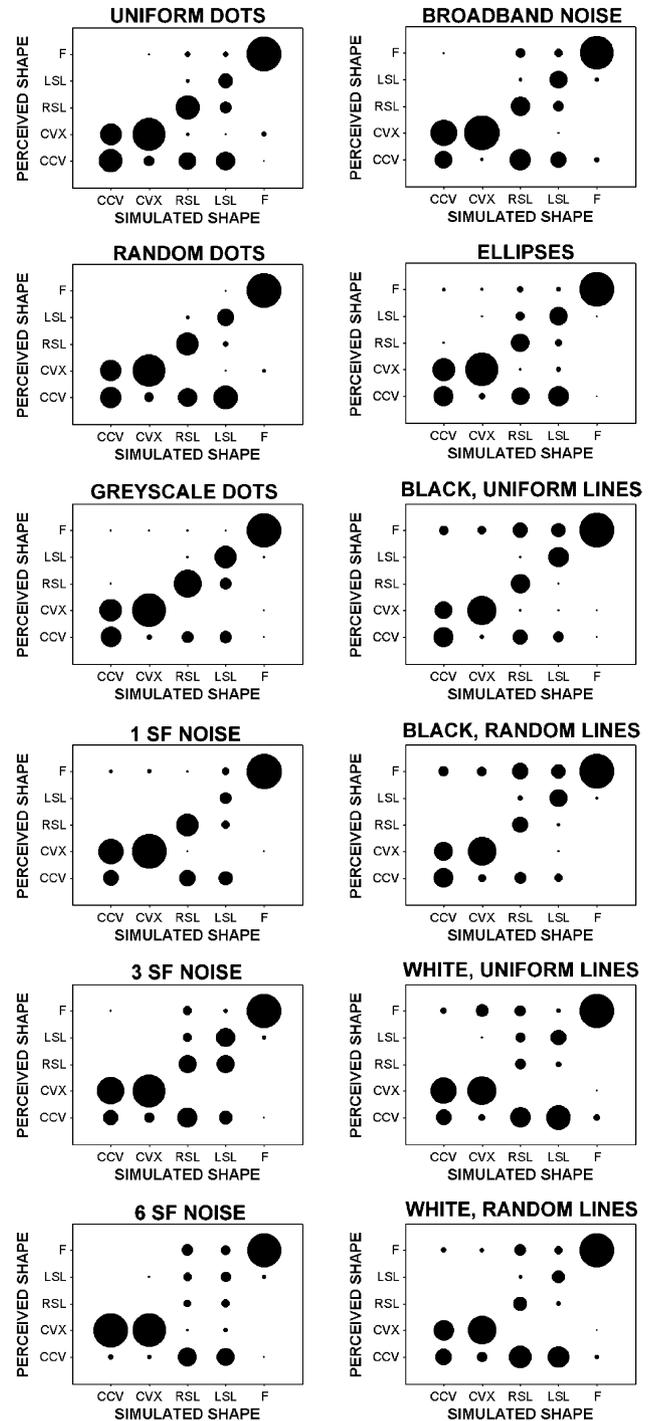


Fig. 11. Averaged data for surfaces pitched backward (Fig. 8). Same format as data in Fig. 7. Observers tend to use the frequency strategy used for upright surfaces, or they identify the shapes correctly.

fact that the frequency along the vertical axis of the image often provided information that was inconsistent with the pitch classification. For example, frequencies along the vertical axis of a concavity pitched forward increase slightly from bottom to top because of the forward pitch, however observers judged the surface to be pitched backward.

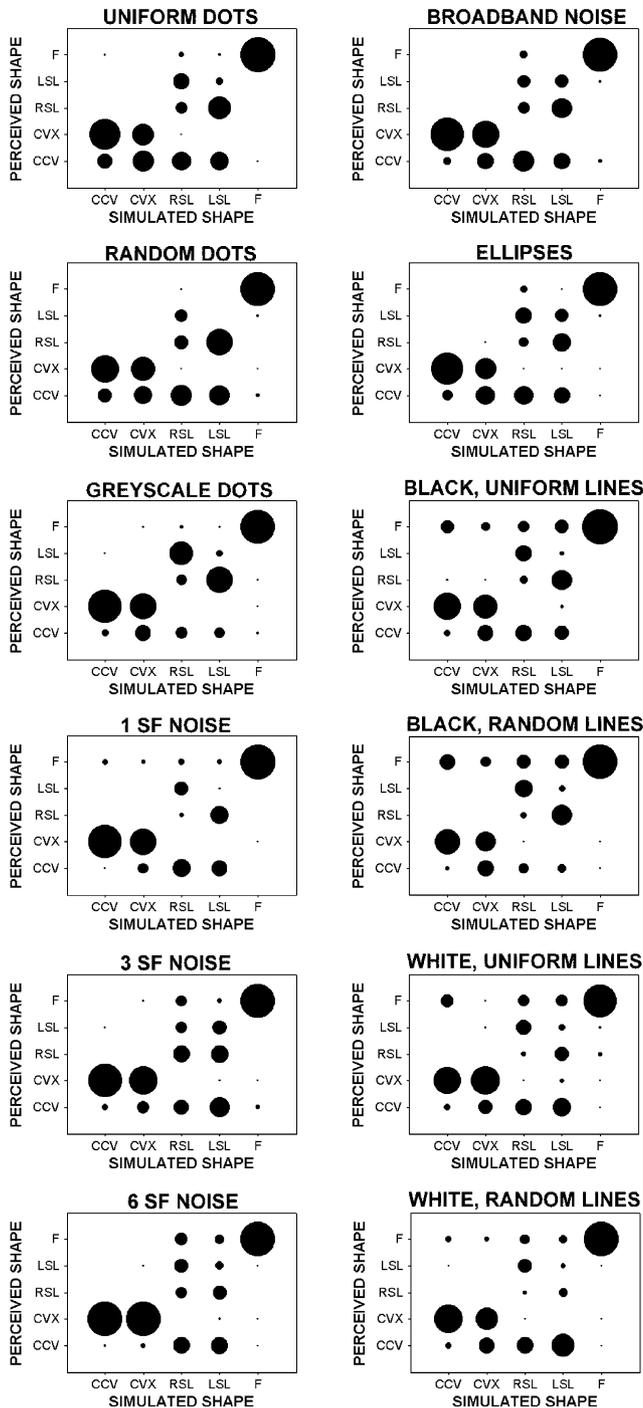


Fig. 12. Averaged data for surfaces pitched forward (Fig. 9). Observers tend to use the frequency strategy or they identify specific oriented flow patterns with the same shapes as they did for surfaces pitched backward, so that responses to right and left slants are now reversed. Downwardly bowed flows were classified as concavities, upwardly bowed flows as convexities, positively oblique as right slants and negatively oblique as left slants.

Although shape and pitch judgments were not made concurrently in the same experiment, the data from Experiment 2 are consistent with the bias found in Experiment 1 in which upwardly and downwardly bowed

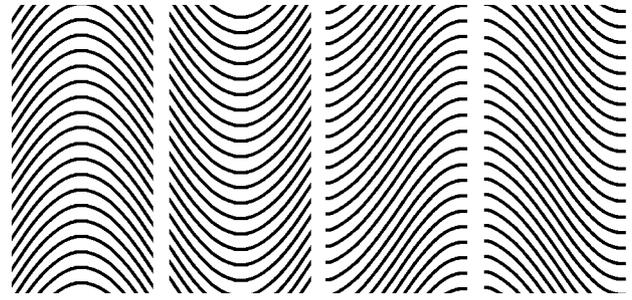


Fig. 13. Oriented contour patterns from Fig. 10 without frequency modulations or perspective changes. The contours within each panel are identical and equally spaced. The left two panels are vertical reflections of one another, as are the right two panels.

flows are perceived as convexities pitched backward and forward respectively.

5. Experiment 3: surfaces containing concave and convex curvatures

The shape percepts thus far have been elicited by perspective images of half-cycles of a corrugated surface. We were interested in determining whether percepts of surface shapes are influenced by surrounding surface shapes. In addition, we were interested to see if the patterns of frequency modulations would still be interpreted as distance in the presence of surrounding curvatures. For example, for an upright corrugation, images of isolated concavities and convexities have similar patterns of frequency modulations, but the image of a central concavity flanked on the left and the right by convexities contains relative frequency information between the center of the concavity and the centers of the convexities that might be sufficient to identify the shapes correctly. In this experiment, perceived shape of the same half-cycles of corrugation in Experiments 1 and 2 was measured using a local depth task when 1.5 cycles of the corrugation were visible.

5.1. Stimuli

We used the uniform polka-dot pattern from Experiments 1 and 2 overlaid onto a surface with the same corrugation parameters. The surface was projected in perspective into larger images spanning $13.5^\circ \times 13.5^\circ$ that were viewed at the distance of 1 m. Viewing these images was identical to viewing the same polka-dotted surfaces in Experiments 1 and 2 through a larger aperture. Through this larger aperture, 1.5 cycles of the corrugation were visible, i.e. three times the number of cycles visible previously. The corrugation was presented in the same four central phases so that the central $9^\circ \times 4.5^\circ$ of the image spanned the same four half-cycle shapes used previously. The surface was also presented

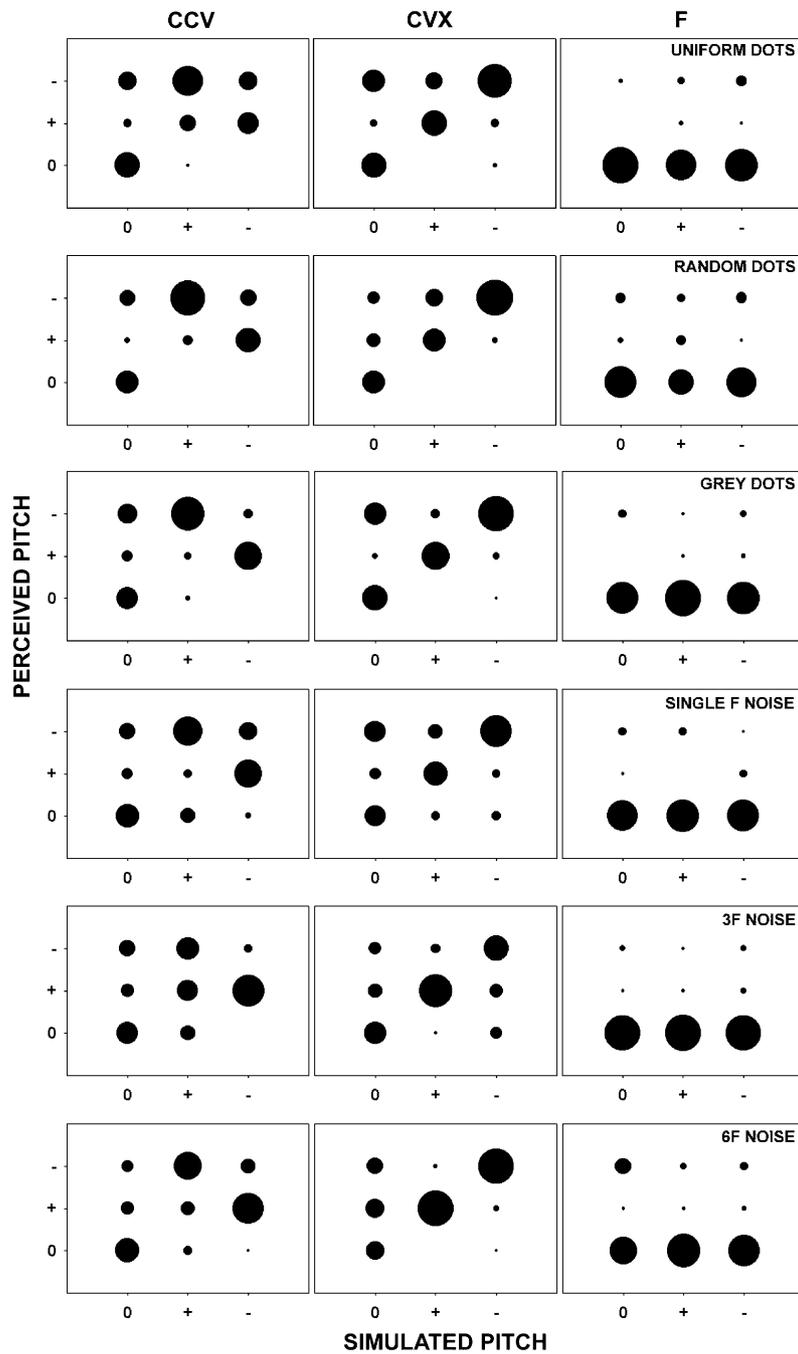


Fig. 14. Averaged data for Experiment 2. Each panel plots for a single pattern and for a single surface shape (left column: concavity, middle: convexity, right: flat) the frequency with which each simulated pitch was reported as each perceived pitch. Correct classifications would result in large dots only along the positive diagonal. For all patterns, flat surfaces were classified as upright regardless of the simulated pitch. For concavities and convexities, observers tended to classify downwardly bowed flows as surfaces pitched forward and upwardly as surfaces pitched backward.

at the same three pitches (0° , $+30^\circ$, or -30°) as Experiments 1 and 2. Fig. 15 shows an example of a stimulus with a central concavity pitched forward.

5.2. Methods

The perceived shape of the central $9^\circ \times 4.5^\circ$ strip of each image (delineated in Fig. 15 by the grey rectangle

which was not visible in the experiment) was measured using a local depth task in which observers judged the relative depth of three horizontally aligned test points in the image. The test points were presented at one of three different vertical heights in the image (all of which are shown in Fig. 15): center, 4.5° above center or 4.5° below center. One of the three dots was always positioned along the vertical mid-line of the image, the other two

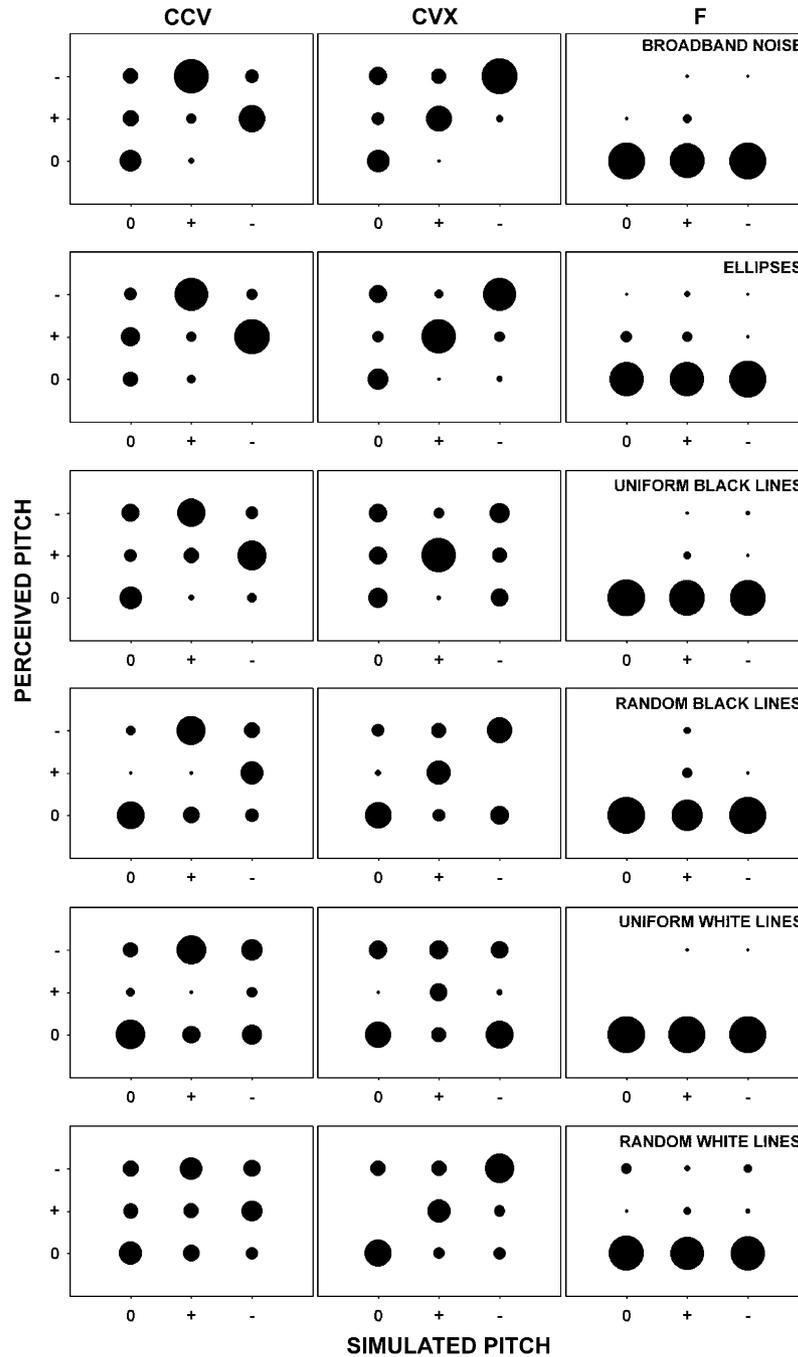


Fig. 14 (continued)

were 2.25° to the left and right of center. In the experiment, the test dots were substantially smaller than the dots in the texture pattern to minimize relative size judgments, and they were colored red for visibility.

In a 4AFC task, observers were asked to indicate the relative depth of the three dots as if they were pasted on the surface by using the toggles on the response box. Observers were told that sometimes the left and right points might not appear symmetrically closer or farther from the central point, but that as long as they both appeared closer or both appeared farther than the central point, they

should respectively respond “concave” and “convex”. They were also told that the surface might sometimes appear upright or pitched towards or pitched away from them, but that they should try to ignore the pitch and focus on judging the relative depth of the three points.

The experiment was divided into three sessions. Each session contained images of the surface at the four different central phases at a single pitch. For each image, the test points appeared at one of the three different vertical heights in the image, for a total of 12 different image/test point combinations. Each of the 12

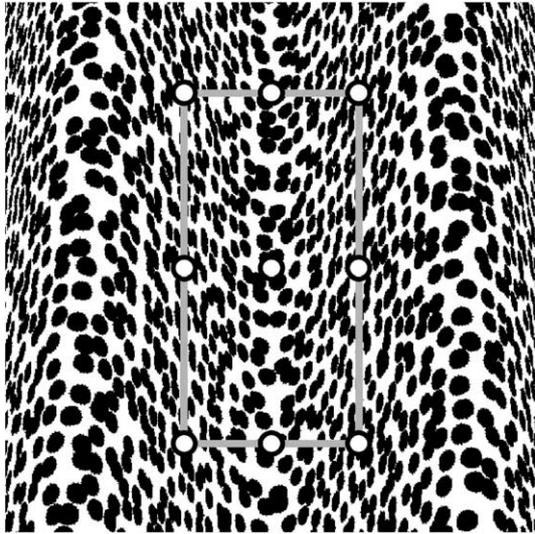


Fig. 15. Example of a stimulus used in Experiment 3. The image contains 1.5 cycles of the corrugation with a central concavity pitched forward. Dimensions of the single-curvature images used in Experiments 1 and 2 are delineated by the grey rectangle that was not present in the experiment. To probe the same surface shapes as those used in Experiments 1 and 2, triplets of test dots were presented at one of three image heights: eye level, 4.5° above, or 4.5° below. In the experiment, test dots were smaller than the dots in the surface texture and red for visibility.

different combinations was presented 10 times for a total of 120 trials per session. The session began with a minute of initial adaptation to a uniform screen of mid-grey at 25 cd/m^2 . Each image was presented for an unlimited time until the observer made a response. Each session lasted approximately 15 min.

Viewing conditions were identical to those used in Experiments 1 and 2. Six of the seven observers from Experiment 1 were run in this experiment.

5.3. Results

Fig. 16 shows results averaged across the six observers. Each panel plots the frequency of perceived shape for each simulated shape. Data for the surfaces presented at the three different pitches are presented in the three different columns—upright on the left, pitched forward in the middle, and pitched backward on the right. Measurements taken from the three test point heights in the image are plotted in the first three rows—the topmost height in the first row, the central height in the second row, and the bottom height in the third row. Data from the same six observers for the uniform polka-dot pattern from the global shape task in Experiment 1 are plotted for comparison in the fourth row. Note that the data from Experiment 1 are for a 5AFC task (including flat surfaces) while the data for this experiment are for a 4AFC task. For both experiments, each simulated shape was presented for the same number of trials, and the dots in each data panel for both experiments have been nor-

malized so that equal areas represent equal percentages of this number of trials. Thus the sizes of the dots for the 4AFC experiment can be directly compared to those for the 5AFC experiment. Note that in the 5AFC experiment (fourth row), observers rarely chose the flat surface as a response for the half-cycle shapes. It is likely that if the 5AFC experiment were run as a 4AFC experiment (omitting the flat fronto-parallel stimuli), the distributions of responses would change very little.

For all three pitches, the data for the three different image heights are similar. This suggests that the different perspective information at the different image heights for the pitched surfaces had little or no effect on perceived shape. At all three pitches, the data similarly show evidence for the frequency strategy under which concavities and convexities were largely classified as convex, and both slants were largely classified as concave. Approximately half of the trials reflect this strategy. Also for all three pitches, there are a proportion of accurate responses as evidenced by the large dots along the diagonal of each panel. This is to be contrasted with the results from Experiment 1 (bottom middle and right panels) for which the oriented flow strategy results in a reversal of percepts when the surface is pitched forward as indicated by the large dots along two anti-diagonals in the bottom middle plot.

These results indicate that observers were able to identify the four shapes correctly more often in the presence of surrounding curvatures than when presented in isolation. More importantly, they indicate that observers are no longer relating particular oriented flows with particular surface shapes at particular pitches. When the surfaces are pitched forward, the oriented flow patterns reverse within the two curvatures and within the two slants, yet observers correctly identify the shapes regardless of the reversals. It is likely that the additional information provided by the surrounding curvatures enables observers to make these correct shape identifications. The increased aperture provides relative frequency information between concavities and convexities and a slightly larger range of frequency changes along the vertical dimensions of the image due to surface pitch. It also provides an increased amount of perspective. Despite the increased number of correct shape classifications however, on approximately half the trials, observers still made incorrect identifications that were consistent with the distance-based frequency strategy. As such, we conclude that the addition of surrounding surface curvatures slightly increases the reports of veridical 3-D shape, but isotropic textures still convey a large number of non-veridical percepts.

6. Discussion

This study identifies conditions under which the shapes of developable surfaces covered with isotropic

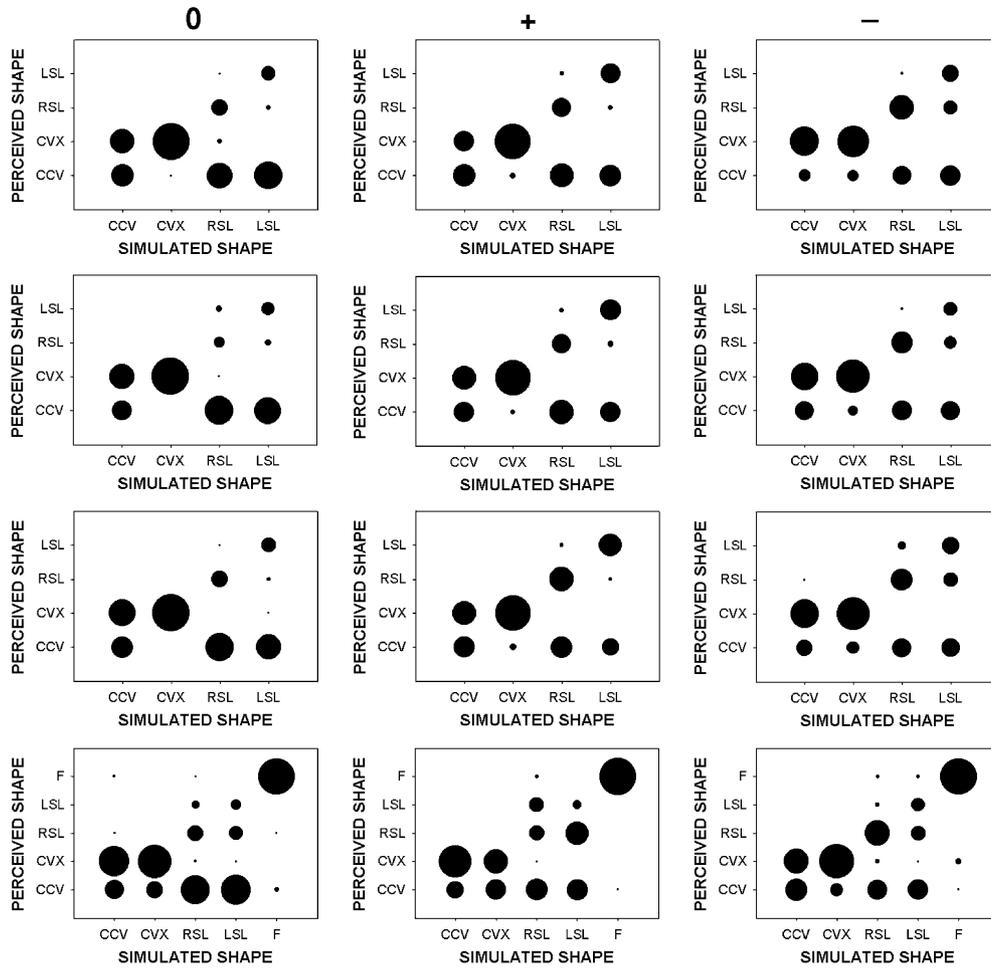


Fig. 16. Averaged data for Experiment 3. Each panel plots perceived vs. simulated shape in the same format as the plots in Fig. 7. The top row shows data for test dots 4.5° above the horizontal mid-line, the second row for test dots along the mid-line, and the third row for test dots 4.5° below the mid-line. Data across the three different image heights are nearly identical. Data from the 5AFC task in Experiment 1 are plotted in the bottom row for comparison. When multiple curvatures are visible, observers tend to classify right and left slants correctly more often when the surface is pitched forward (middle column, top three rows) compared to when only single surface curvatures are visible (middle column, bottom row).

textures are misperceived and conditions under which they are perceived veridically. Both the correct and incorrect shape classifications can be accounted for by the same perceptual strategies.

Images of upright developable surfaces contain frequency modulations along the axis of maximum surface curvature. The data for upright surfaces are consistent with a strategy in which these one-dimensional modulations are attributed to changes in distance rather than slant such that low frequencies indicate closer portions of the surface rather than fronto-parallel portions and high frequencies farther rather than slanted portions. The data are consistent with this strategy despite the fact that the one-dimensional frequency modulations along the axis of maximum curvature are physically consistent only with changes in slant. When frequency modulations are misinterpreted as changes in distance, observers misclassify concave half-cycles of a sinusoidally corrugated surface as convex, and slanted half-cycles as concave.

Images of pitched surfaces exhibit similar patterns of frequency modulations as the images of upright surfaces, however in addition they contain oriented flows along projected lines of maximum surface curvature. The data for pitched surfaces are consistent with two concurrent strategies. The first is the same frequency strategy used for the upright surfaces, resulting in the misclassification of concavities as convex and slants as concave. The second strategy is one in which each of these four patterns of oriented flows is related to each of four surface shapes—upwardly bowed flows as convexities, downwardly bowed flows as concavities, positively oblique as right slants and negatively oblique as left slants. Each flow pattern depends not only on the shape of the surface, but also on its pitch; patterns for concavities and convexities pitched backward are reversed when the surfaces are pitched forward, and similarly for left and right slants. For many patterns, a large proportion of surface shapes were judged correctly when they were pitched backward. However, when the same surfaces were pitched forward,

the shape judgments were reversed, consistent with the reversed oriented flow patterns: concavities and convexities were confused, as were right and left slants. The fact that shapes were classified correctly more often when they were pitched backward rather than when they were pitched forward (especially for right and left slants) appears to contradict previous studies suggesting the existence of a bias towards perceiving the ground plane (Gibson, 1950). The data also appear to reveal a bias in which both upwardly and downwardly bowed flows are perceived as convex surfaces.

The results from Experiment 2 suggest that observers relate particular oriented flows with surface shapes at particular pitches, so that upwardly bowed flows indicate a convexity that is pitched backward, and downwardly bowed flows indicate a convexity pitched forward. Surface pitch was classified correctly for convexities more often than for concavities or flat surfaces. This result was revealed through a pattern of paired pitch classifications: a concavity pitched forward and a convexity pitched backward both exhibit upwardly bowed flows and were classified as pitched backward, while a concavity pitched backward and a convexity pitched forward both exhibit downwardly bowed flows and were classified as pitched forward. These results are thus consistent with the bias of perceiving upwardly and downwardly bowed flows as convex surfaces.

Experiment 2 also showed that, for the present stimulus conditions, observers tend to classify isotropically textured shapes pitched forward or backward as upright. While it has been previously shown that observers tend to underestimate surface slant from texture cues (Gruber & Clark, 1956; Turner et al., 1991), we know of few studies showing this underestimation for surfaces that are pitched backward (besides Gibson, 1950). Ideal observer analyses have shown that the reliability of texture information increases with surface slant (Knill, 1998c). It is possible that surfaces at steeper pitches than those used here would not have been perceived as upright.

The bias towards surfaces pitched backward cannot be the result of differential amounts of information in the images of surfaces pitched in one direction as opposed to those pitched in the other direction. As Fig. 10 shows, the orientation modulations consistent with the four surfaces pitched forward are exact vertical mirror images of the same surfaces pitched back. Observers instead appear to be relating particular patterns of orientation modulations with particular surface shapes at particular pitches. The relation is made also despite inconsistent changes in spatial frequency. In addition to frequency modulations occurring along the axis of maximum curvature, the images also contain slight changes in frequency along the axis of minimum curvature depending on the pitch of the surface. For surfaces pitched forward, the frequency increases slightly from the bottom to the top of the image, and for surfaces pitched backward, the

frequency decreases slightly from bottom to top. Observers appear to be relating each pattern of orientation flow with a particular shape at a particular pitch regardless of whether the frequencies increase or decrease from top to bottom along the vertical axis. For the pitches used in this study, these vertical frequency modulations apparently do not play a critical part in observers' shape or pitch judgments. Consistent with this, results from Experiment 2 and previous studies (Cutting & Millard, 1984; Knill, 1998b) have shown that changes in spatial frequency or size/density along this axis are not effective for conveying the pitch of flat surfaces.

It is important to note that while some proportion of shapes were correctly classified for surfaces pitched backward, a large proportion of them were also incorrectly classified at this pitch consistent with the frequency strategy, even in the presence of oriented flows. Our previous work (Li & Zaidi, 2001a; Zaidi & Li, 2002) has indicated that the only patterns of orientation modulations that can clearly disambiguate between concavities, convexities, right slants, and left slants are patterns of contours or flows shown in Fig. 17: inwardly bowed flows for concavities, outwardly for convexities, radially outward from right to left for right slants, and radially outward from left to right for left slant. These critical patterns arise from particular oriented components of the surface texture, the orientations of which depend on the pitch of the surface (Zaidi & Li, 2002). For example, for the upright surfaces in Fig. 17, they arise from the horizontal component. If the same surfaces were pitched forward by 30°, these same critical patterns would arise from components oriented $\pm 30^\circ$ from the horizontal. In the absence of the critical patterns, observers largely misclassified concavities as convex. These critical patterns were not visible in any of the isotropic stimuli used in the present study. Although isotropic textures do contain components at all orientations, it remains an empirical question as to why the critical flows are not visible. It is possible that the contrasts of the critical components are masked by components at neighboring orientations. Oriented flows of the isotropic textures are visible only when the surfaces are pitched, and these

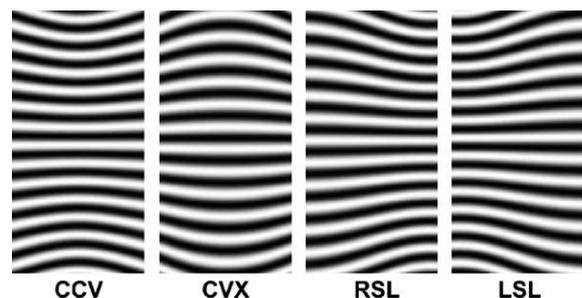


Fig. 17. Critical patterns of orientation modulations required for distinguishing concavities, convexities, right and left slants, regardless of surface pitch.

flows follow projected lines of maximum curvature, which our work has shown to be sufficient for conveying veridical shapes of upright surfaces (as in Fig. 17), but insufficient for pitched surfaces (Zaidi & Li, 2002).

While Experiments 1 and 2 utilized isolated half-cycle surface curvatures, the results from Experiment 3 showed that if surrounding concave and convex curvatures are simultaneously visible, some shape judgments improve. The visibility of surrounding curvatures provides relative frequency information between concavities and convexities so that while local frequencies are still low at both curvature extrema, the frequencies at a convexity are lower than those at a concavity. For pitched surfaces, the addition of surrounding curvatures also provides the juxtaposition of surface curvatures all converging in the same direction. In addition, the increased vertical extent of the projected image provides increased perspective information. In the presence of this additional information, observers' judgments of convexities and concavities improved, however concavities were still misperceived as convexities as often as they were perceived correctly. In addition, slants were still misperceived as concavities as often as they were in Experiments 1 and 2, which suggests that frequency modulations are still being misinterpreted as distance. For surfaces that are pitched forward, the number of correct slant classifications increased so that observers no longer appear to be relating particular flow patterns with particular surfaces pitched backward. That is, the presence of surrounding curvatures appears to reduce the bias for perceiving surfaces pitched backward.

To summarize, our results show that frequency and orientation modulations are used in different strategies to determine surface shape and surface pitch, and that these same strategies account for veridical and non-veridical percepts of developable surfaces. One-dimensional frequency modulations in the perspective image arising from surface slant are misinterpreted as surface distance, resulting in the misperception of concavities as convexities and slants as concavities. Orientation modulations in the image are used in two different ways: in some cases, patterns of orientation modulations are interpreted, often incorrectly, as particular surface shapes pitched backward. In other cases, bowed flow patterns are specifically perceived as convex surfaces, thereby additionally increasing the frequency with which concavities are reported as convex. The presence of surrounding surface curvatures only slightly improves the veridicality of the percepts.

Acknowledgements

The authors wish to thank Seema Patel for her help with Experiments 2 and 3, and A. Fuzz Griffiths and Rocco Robilotto for comments on the manuscript. This

work was supported by NEI grant EY13312 to Qasim Zaidi, and was presented in part at the Second Annual Visual Sciences Society Meeting in Sarasota, Fla., May, 2002, and the European Conference on Visual Perception in Glasgow, Scotland, August, 2002.

Appendix A

If a grating of spatial frequency F and orientation ω is overlaid onto a surface slanted about a vertical axis at θ degrees away from the fronto-parallel plane and projected in perspective so that the surface (and the image plane) are viewed at a distance of D , the projected spatial frequency F' along the line of sight has been derived in Li and Zaidi (2001a) to be

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

where $\omega_p = \omega + \pi/2$. The steepest slants in the sinusoidal half-cycle surface shapes used in our experiments were approximately 71° from the fronto-parallel plane. For a vertical grating ($\omega = \pi/2$) overlaid onto a surface slanted at 71° and viewed at the viewing distance of $D = 100$ cm, an arbitrary frequency F will be increased by a factor of 6 from its original value. An equivalent increase in F could be obtained if a fronto-parallel surface is viewed at a distance of 300 cm. The amplitude of the sinusoidally corrugated surfaces used in our experiments was 7 cm so that the distance between the observer and the surface only varied between 93 and 107 cm. The frequency modulations in the image can thus be largely attributed to the surface slant and not the changes in distance.

References

- 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—General*, 113(2), 198–216.
- Gibson, J. J. (1950). The perception of visual surfaces. *American Journal of Psychology*, 63, 367–384.
- Gruber, H. E., & Clark, W. C. (1956). Perception of slanted surface. *Perceptual and Motor Skills*, 6, 97–106.
- Hel Or, Y., & Zucker, S. W. (1989). Texture fields and texture flows. *Spatial Vision*, 4(2–3), 131–139.
- Jamar, J. H., Campagne, J. C., & Koenderink, J. J. (1982). Detectability of amplitude- and frequency-modulation of suprathreshold sine-wave gratings. *Vision Research*, 22(3), 407–416.
- Knill, D. C. (1998a). Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Research*, 38(11), 683–1711.
- Knill, D. C. (1998b). Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture. *Vision Research*, 38(17), 2635–2656.
- Knill, D. C. (1998c). Surface orientation from texture: ideal observers, generic observers and the information content of texture cues. *Vision Research*, 38(11), 1655–1682.

- Knill, D. C. (2001). Contour into texture: information content of surface contours and texture flow. *Journal of the Optical Society of America A*, 18(1), 12–35.
- Li, A., & Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research*, 40(2), 217–242.
- Li, A., & Zaidi, Q. (2001a). Erratum to Information limitations in perception of shape from texture. *Vision Research*, 41(22), 2927–2942 (*Vision Research* 41 (2001) 1519–1534).
- Li, A., & Zaidi, Q. (2001b). Information limitations in perception of shape from texture. *Vision Research*, 41(12), 1519–1533.
- Turner, M. R., Gerstein, G. L., & Bajcsy, R. (1991). Underestimation of visual texture slant by human observers: a model. *Biological Cybernetics*, 65(4), 215–226.
- Zaidi, Q., & Li, A. (2002). Limitations on shape information provided by texture cues. *Vision Research*, 42(7), 815–835.