Investigating the Effect of Texture on 3D Shape Perception

Sunghee Kim

Computer Science and Engineering
University of Minnesota
email: skim@cs.umn.edu

August 29, 2002

A report to the committee for the
Ph.D. Preliminary Oral Examination

1. INTRODUCTION

A key objective in the field of visualization is to design and implement algorithms to effectively convey scientific data through images so that the essential features of the data can be easily understood. It is important that the projected images facilitate the accurate and intuitive understanding of the data. When we choose to render a surface or object, we have tremendous latitude in determining how we want to model its material properties. The most common practice is to use a simple Phong shading model without any surface texture, because it is simple to implement and is the default on most systems. However this model may not be optimal for all purposes, and in particular is not optimal for shape representation. Unfortunately, the existing theories on shape perception do not provide sufficient guidance to definitively answer the question of how best to define the surface material properties of an object in order to best facilitate the accurate understanding of its shape.

The overall objective in this report is to investigate the effects of texture pattern characteristics on the perception of the three-dimensional shape of smoothly curving surfaces that have the potential to yield fundamental theoretical insights into what works, what does not, and more importantly *why*. The answers to these questions have important potential impact across a wide range of visualization applications, from molecular modeling to radiation therapy treatment planning, in which scientists need to attain an accurate understanding of the shape of smooth, complicated, and arbitrarily curving surfaces in their data.

Because of the limitations of classical texture mapping software and algorithms, nearly all studies investigating the effect of surface texture on shape perception have been restricted either to the use of developable surfaces (which can be unrolled to lay out flat in the plane)or to the use of procedurally defined solid texture patterns, whose characteristics are in general independent of the geometry of the surfaces to which they are applied. For several years we have believed that important new insights into the effects of texture on shape perception might be gained through studies conducted under less restrictive surface and texture pattern conditions.

2. PREVIOUS WORK

Although artists have been using texture characteristics as a pictorial cue for a long time, scientific study of texture as a cue in visual perception started in 1950 with Gibson [7]. Gibson hypothesized that the perception of the surface slant at a point depends on the rate of change of the texture density – projective distortion of surface texture. In his experiment, he used a brick-like regular texture that consisted of evenly distributed rectangular elements with linear perspective and an irregular texture with a relatively uneven distribution. Observers were asked to judge the slant of the surface by viewing textures that were slanted in varying degrees. His results confirmed that the gradient of texture density is a determinant of apparent optical slant.

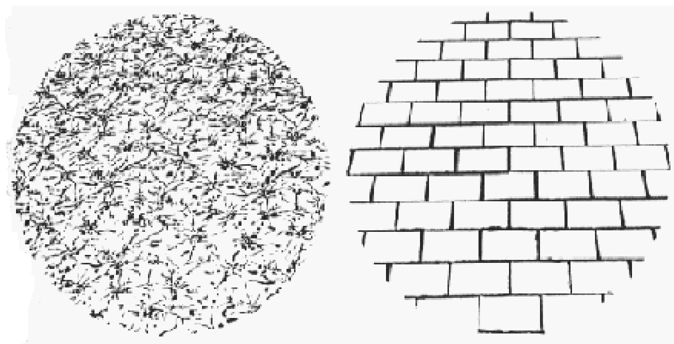

Figure 1. Textures used in Gibson's experiment. Left: irregular texture, Right: regular texture.

In the computer vision community, two models of shape-from-texture have been proposed. The first model assumes that the surface texture is *homogeneous*, i.e., the observer can assume that the density of the texture is nearly constant (homogeneous) on the surface and use the variation of the density in the image to judge the shape and orientation of the surface. The hypothesis of this model is that the observer extracts three-dimensional shape information by comparing the change in texture gradient from one patch to another. This assumption can be applied to other characteristics of the texture pattern such as the area of the texture elements. The human visual system can then use the variation of one or more texture characteristics to determine surface shape and orientation. In [22], Marinos and Blake show how a maximum likelihood estimator can be constructed for homogeneous texture under full perspective transformation.

The second model assumes that the surface texture is *isotropic*, i.e., there is no dominant direction in the texture. The hypothesis of this model is that the observer can infer the shape by measuring the deviation of the orientation distribution from isotropic. This approach was first proposed by Witkin [29]. Witkin noted that the *foreshortening* effect (the fact the image of a slanted pattern is systematically compressed in the direction of slant) can be used as a cue to surface orientation. While the texture gradients suggested by Gibson vanish under orthographic projection, the foreshortening effect can still be observed under orthographic projection. Blake and Marinos [1] proposed an iterative algorithm for computing the maximum likelihood estimator. Gårding [6] extended Witkin's theory to perspective projection and presented an efficient algorithm for estimation of surface orientation.

Rosenholtz and Malik [24] examined the above two models for human perception of shape-from-texture. In their study, observers viewed planar surfaces for a number of different orientations, and for both isotropic and anisotropic textures. Observers indicated the perceived orientation of the plane by adjusting a surface attitude probe in the center of the image modeled after that used in [16 Koenderink *et al.*]. The results suggest that the human visual system uses both homogeneity and isotropy as cues for determining surface shape and orientation.

Although the mechanisms of texture's effect on shape perception are not yet completely understood, numerous studies over the recent years have found evidence that the accuracy of observers' judgments of surface orientation and curvature can be affected significantly by the presence of a surface texture pattern. However, attempting to use texture to facilitate veridical shape perception can be difficult. We must avoid employing texture in a way that masks shape information or leads to an exaggerated or inaccurate perception of surface curvature or orientation, as occurs in op art for example.

Ferwerda *et al.* [5] showed how texture can be used to *mask* surface shape features. They developed a computation model of visual masking that predicts how changes in the contrast, spatial frequency, and orientation of the surface texture pattern or changes in the tessellation of the surface will affect the visual masking.

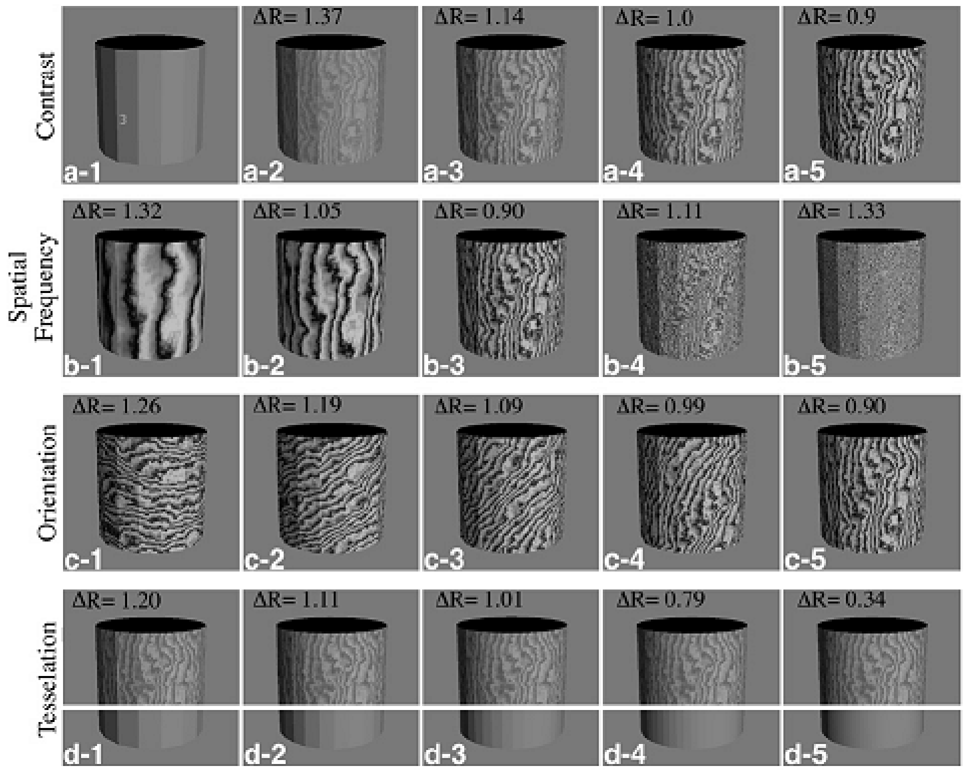

Figure 2. Examples of visual masking by Ferwerda *et al.* [5] Top three rows show how change in surface texture pattern alters the masking effect, Bottom: how change in surface tessellation affects the masking.

Cumming *et al.*[4] examined whether the human visual system uses a global or local solution for shape-from-texture. They generated stimuli by texturing horizontal circular cylinders with patterns of ellipsoids (with horizontal major axis) before projection. This ensures that the ellipticity of any single texture element cannot be used as a gauge of surface orientation. The compression, density, and area of texture elements were introduced and modified independently throughout the experiment. An isotropic texture with randomly oriented ellipsoids was included as shown in Figure 3. Observers judged whether cylinders appeared elongated or flattened relative to a circular cylinder. Their results support that compression gradient cue was more effective when the texture was isotropic than when it was anisotropic. They suggest that texture cue provided by change in texture element compression is most reliable in conveying shape and that surfaces textured with anisotropic texture can impede shape perception.

Recent findings support the idea that the facility with which we can accurately perceive surface shape in the presence of texture depends not only upon the intrinsic characteristics of the texture pattern itself but also upon how the pattern is laid down over the surface [27 Todd *et al*., 26 Stevens, 21 Mamassian *et al.*, 18 Li *et al.*, 19 Li *et al.*, 20 Zaidi *et al.,* 14 Knill].

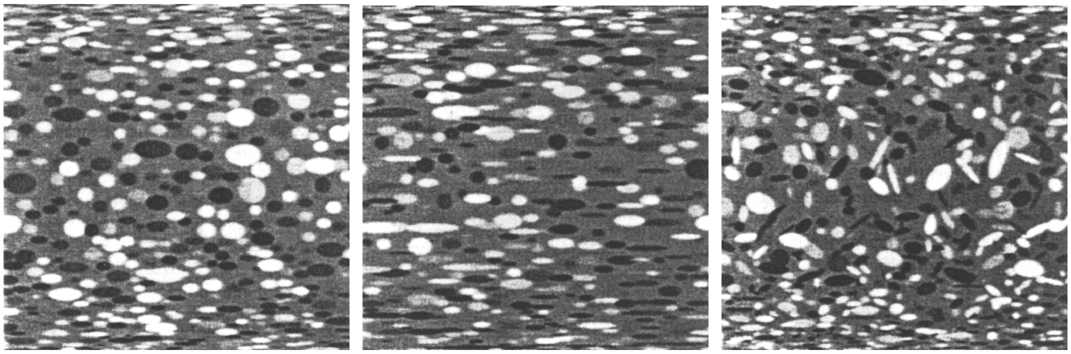

Figure 3. Example stimuli used in Cumming *et al.*'s study. Texture elements have Left: maximum elongation of 1.5; Center: maximum elongation of 3.0; Right: maximum elongation of 3.0, random orientation

Todd and Reichel [27] suggest that texture composed of short or long, straight or curved, continuous or disjointed contours can give strong 3D shape percepts. They argue that shape perception will deteriorate if the length of the contours is too short to show their relative orientations, or if the pattern masks changes in average orientation. The results show that the human visual system determines 3D surface structure from the statistical distributions of contour orientation within each local neighborhood.

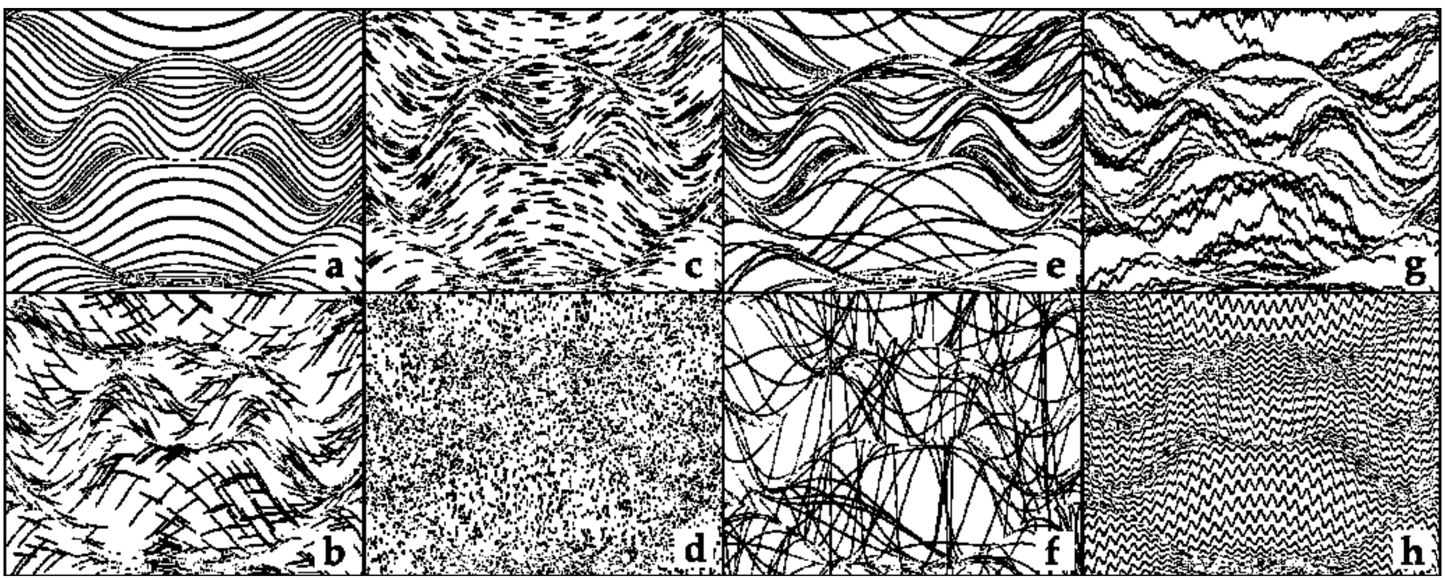

Figure 4. Examples of a curved surface textured with different patterns. Notice that texture patterns in (a) ~ (c), (e) (g) give strong shape percepts while patterns in (d), (f), (h) do not.

Stevens [26] showed that surface orientation at a point can be computed based on the relationship between the geometry of the surface and the surface markings. Gaussian curvature at any surface point is the product of the first and second principal directions (or maximum and minimum curvatures) at that point. The surface normal direction vector can be calculated by finding the vector that is perpendicular to the plane in which the first and second principal directions lie. Conversely, if the maximum and minimum curvatures at a surface point are given, the surface normal at that point can be estimated fairly accurately. He proposed that when we see line markings on a surface, we assume the lines follow the principal directions.

Mamassian and Landy [21] investigated the perception of solid shape from line drawings. Observers viewed one of the three Figures (in Figure 5) in random order and asked to report whether their first impression of the shape was that of a saddle-shaped (hyperbolic) surface or of an egg-shaped (elliptic) surface. They reported that even though the surface contour information provided by simple line drawings is quite ambiguous, the performance of shape perception was stable. They introduced three observer preference assumptions: 1) a preference for a convex surface; 2) a preference for surface contours aligned with the principal curvature; 3) a preference for a surface orientation consistent with an objected viewed from above.

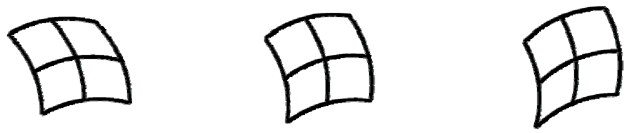

Figure 5. Figures shown to observers in [21]. Left: compressed along the axis of symmetry, Right: stretched along the axis of symmetry, Center: an intermediate transition between the left and right line drawings.

Li and Zaidi [18, 19] examined the ability of human observers to perceive whether a textured surface is convex or concave. Their stimuli consisted of surfaces that were textured in simple grating and plaid texture patterns, complex plaids containing energy at multiple orientation and frequencies, and filtered noise patterns (Figure 6). These textured surfaces were corrugated sinusoidally in depth about the horizontal axis and projected in perspective. In a series of local depth judgments, observers were asked to indicate which of the two locations on the surface around a central fixation point appeared closer to them in depth. The results support the hypothesis that veridical ordinal depth is seen only when the projected pattern followed lines corresponding to maximum curvature of surface and that lines of minimum curvature can also be important for conveying shape. They infer that since the correlated orientation changes along maximum and minimum surface curvature are lost in orthographic projection, texture is a useful cue in shape perception in perspective projection.

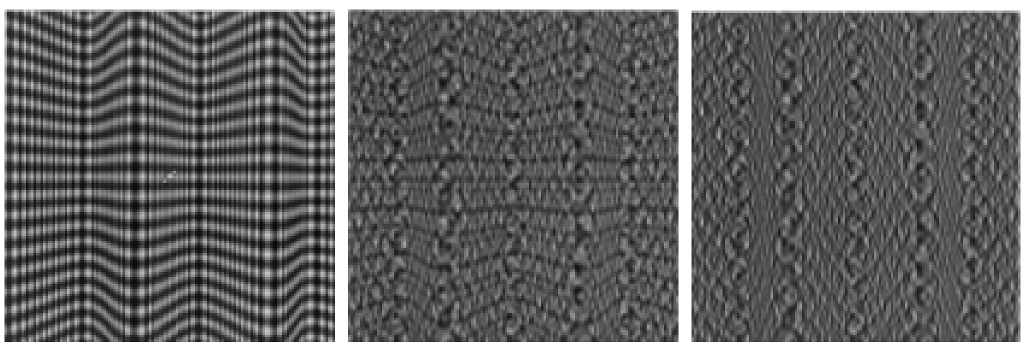

Figure 6. Example stimuli shown to observers. Left: horizontal vertical plaid pattern, Center: octotropic plaid pattern, Right: octotropic plaid pattern without the component parallel to the axis of maximum curvature. Notice that the pattern on the right does not convey the shape of the surface as well as the patterns on the left and center.

Zaidi and Li [30] further examined the role of orientation of the texture pattern in conveying three-dimensional shape. Their hypothesis was that both perspective projection and texture orientation are necessary for veridical shape percepts of convex and concave surfaces. In their experiment the observers were presented with corrugated surfaces pitched in varying degrees. They argued that "different pairs of geodesics are necessary to convey veridical shape at different pitches of the corrugation, each pair arising from components symmetrically oriented around the axis of maximum curvature." Figure 7 shows simulated 3D corrugations used in their experiment. The first two columns show convex and concave surfaces pitched at 0 degree, next two columns show surfaces pitched at 22.5 degree (tilted backward), and the last two columns show surfaces pitched at –22.5 degrees (tilted forward). Observers correctly perceived shape for all pitch degrees in (a), but reported correct shape only for 0 pitch degree. In (c) they reported concavities as convexities for 0 pitch degree. Observers had difficulty making shape judgments for 0 pitch degree but reported strong shape perception for others in (d). They were able to identify correct shape for all cases in (e).

	texture pattern	convex concave 0 degree · convex concave 22.5 degree · convex concave - 22.5 degree
(a)	octotropic plaid	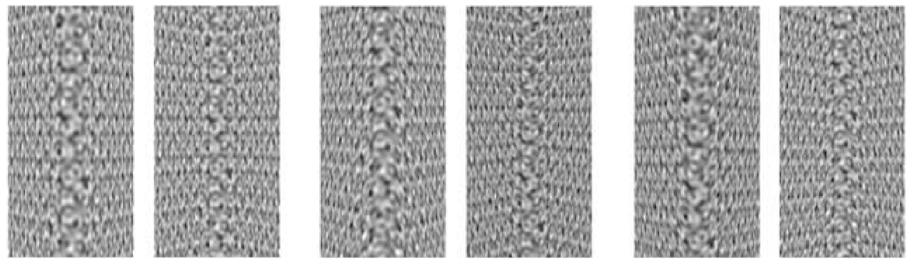
(b)	with horizontal and vertical components only	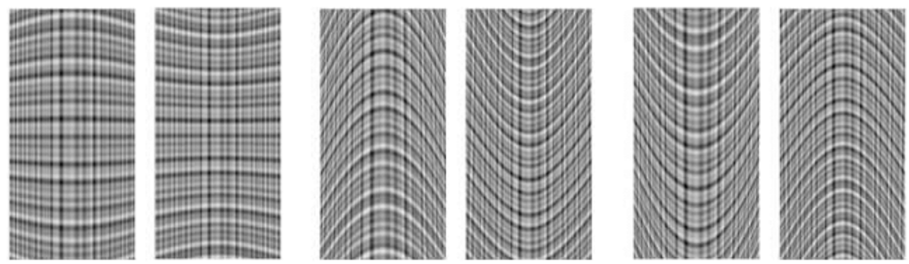
(c)	without horizontal components	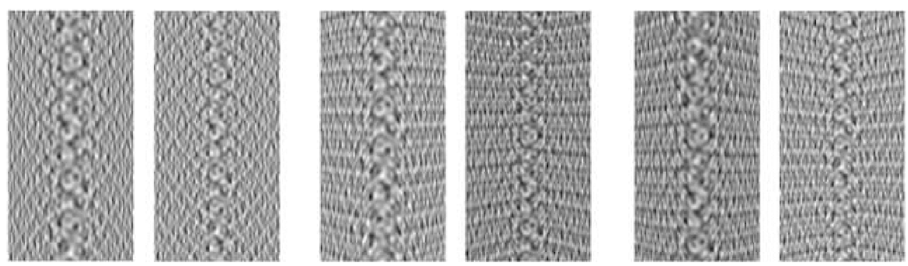
(d)	with diagonal components only	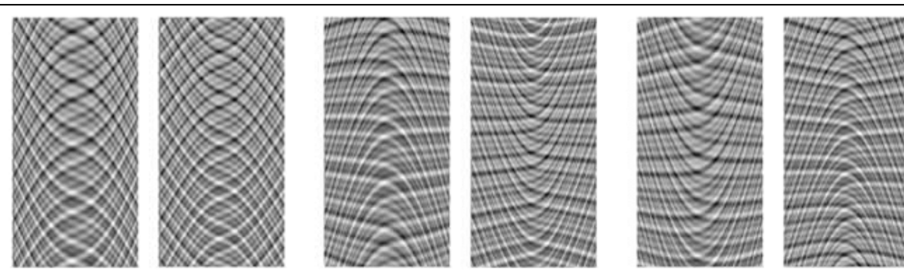
(e)	without diagonal components	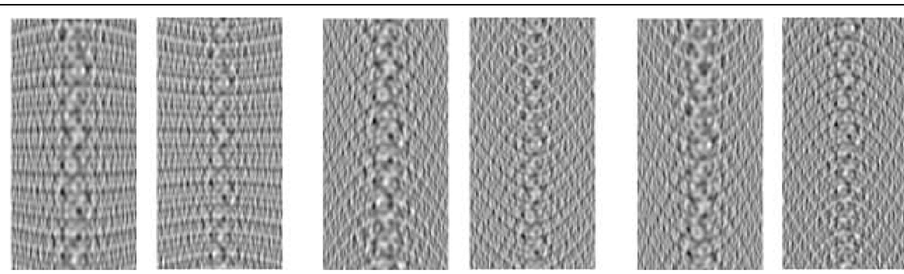

Figure 7. Example stimuli used in Zaidi *et al.* [30] Convex and concave surfaces are textured in patterns of different components.

According to Knill [14], “many of the most perceptually salient texture patterns have a strong flow-like structure, resulting from the directional nature of the surface textures from which they project.” Knill demonstrates that across developable surfaces any homogeneous anisotropic texture pattern will appear to flow along parallel geodesics (curves that do not turn in the surface). He suggests that our visual system uses a shape-from-contour mechanism to infer shape from the systematic projective distortion or flow of the pattern.

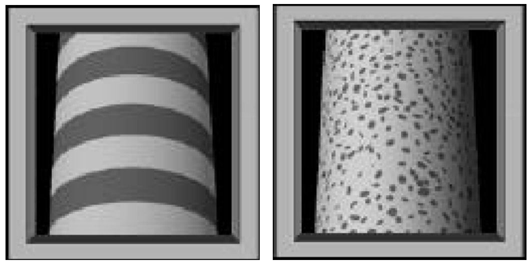

Figure 8. Examples of surface markings that generate strong cues to shape. Left: shape-from-contour, Right: shape-from-texture

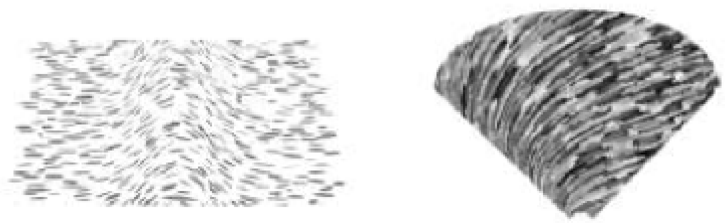

Figure 9. Examples of homogeneous oriented texture patterns on developable surfaces. Left: texture flows along the line of maximum curvature, Right: texture flows along the cone but not the line of maximum curvature.

Inspired by these findings that seem to imply that surface shape may be perceived most accurately from line-like markings when they follow the lines of curvature, we first sought to further experimentally investigate the effect of the direction of surface texture pattern on the accuracy of observers' shape judgments.

The rest of the paper is structured as follows. Sections 3 ~ 5 will describe the three experiments conducted in our lab in recent years; details of stimuli preparation, the tasks of observers, results and analysis will be discussed. After a brief summary, section 7 will present directions and topics for future research.

3. EXPERIMENT 1

3.1 OVERVIEW

In this experiment, we examined the effect of the presence and direction of luminance texture pattern anisotropy on the accuracy of observers' judgments of 3D surface shape. Specifically, we sought to answer the questions: can an anisotropic pattern that follows the principal directions show shape more effectively than a pattern in which the direction of anisotropy follows some other path? Is the pattern that follows the principal direction more effective than an isotropic pattern? Are the effects the same in the case of shaded displacement texture? To what extent are these effects mitigated by stereo viewing?

We measured the accuracy of observers' shape perception judgments by having them manipulate an array of surface attitude probes [16 Koenderink *et al.*] so that the perpendicular extensions appeared to point in the direction of the local surface normal. Stimuli were displayed as flat images in the first series of trials, and then the process was repeated with the same surfaces displayed in stereo. In the flat viewing condition, performance was significantly better in the cases of principal direction texture and isotropic patterns than in the cases of the sinusoidally swirly and uniform patterns. In the stereo viewing condition, accuracy increased for all texture types, but was still marginally greater in the cases of the isotropic and principal direction patterns than under the other anisotropic conditions.

These results are consistent with a hypothesis that texture pattern anisotropy can impede surface shape perception when the elongated markings are oriented in a way that is different from the principal direction.

3.2 STIMULI

The stimuli we used in our experiments were cropped images of the front-facing portions of textured level surfaces rendered in perspective projection. The images were rendered with a hybrid renderer [9 Interrante *et al.*] which uses raycasting [17 Levoy] together with a Marching Cubes algorithm [20 Lorensen] for surface localization. The volumetric test data from which we extract these surfaces is a three-dimensional dose distribution calculated for a radiation therapy treatment plan. We chose to use the radiation data as our test bed, rather than a more restricted type of analytically-defined surface, because this is typical of the kind of data whose shape features we hope to more effectively portray through the use of surface texture.

The first step in image generation was to define the texture patterns that would appear on the level surfaces. We used a high-quality three-dimensional line integral convolution algorithm [25 Stalling] to synthesize solid textures in the vicinity of the selected level surface. Beginning with a three-dimensional array of binary noise, line integral convolution produces an output texture in which the input values are correlated along the directions indicated by an accompanying vector field. We defined four different vector fields to produce four different types of texture patterns.

The procedure that we used to obtain the principal direction vector field is fully described in [10 Interrante]. Briefly, we compute an orthogonal frame at each sample point in the 433x357x325 volumetric dataset. We define a frame vector to be in the direction of the gray-level gradient, which is the normal to the level surface that passes through the sample. Using Gaussian-weighted central differences in the axial directions over the 3x3x3 area surrounding the sample point we compute the gradient. Finally, we estimate the 2nd Fundamental Form [15 Koenderink] from the Gaussian-weighted central differences of the gradients trilinearly interpolated at sample positions over a 3x3x3 grid aligned with the local frame, diagonalize to obtain the 2D principal directions (eigenvectors) and principal curvatures (eigenvalues) in the tangent plane, and convert to 3D object space coordinates. The direction corresponding to the eigenvalue with the greatest unsigned magnitude is saved in the 3D principal direction vector array and used to create the first anisotropic texture ('*pdir'*).

The remaining 3D vector fields are as follows. First, we obtain the vector field of uniform directions ('*udir*') by taking at each point the direction given by the intersection of the tangent plane with the plane orthogonal to the z axis that passes through the sample point: udirx = –ny, udiry = nx, udirz = 0, where (nx, ny, nz) is the surface normal or gradient. We obtain the vector field of random directions ('*rdir'*) that is used to create the isotropic texture pattern by rotating the uniform direction previously obtained at each point by a random angle θ_1 about the surface normal, $-\pi/2 < \theta_1 < \pi/2$. Finally, we obtain the vector field of coherently varying directions ('*sdir*') used to create the anisotropic texture pattern that contains lines with non-zero geodesic curvature by rotating the original uniform direction about the surface normal by an angle $\theta_2 = 10\,\pi\,(x+y+z)/n$ where (x,y,z) is the index of the sample point in the volume and *n* is the total number of sample points in the 3D array.

During rendering, the intensity value interpolated from the 3D texture at the ray/surface intersection point is taken as the base color of the surface at the ray surface intersection point, then Phong shading is applied to obtain the final surface color. We rendered 48 test images for the experiment, 24 for the left eye views and 24 for the right eye views, using the four different textures applied to views from six different vantage points around a single level surface. Figure 10 shows four of these images, all computed for the same viewing position. In order to avoid the potentially confounding influence of shape from contour information, we cropped each image to a 400x400 pixel region that did not contain any silhouette edges of the object.

3.3 EXPERIMENTAL TASK

Our goal was to design an experimental task that could reveal the effect of different texture types on the accuracy and efficiency of an observer's perception of the global 3D shape of a displayed object (shape from a glance). It is well known that our visual system does not build up an estimate of shape from the accumulation of isolated individual local estimates of surface heading. Instead our visual system obtains shape understanding from the comparative relationships between nearby points. Therefore we presented an array of probes [16 Koenderink *et al.*] that completely covered the central area of the presented surface. Each probe had a circular base with a perpendicular extension and observers were asked to adjust each probe by pulling on its handle until the perpendicular extension appeared to point in the surface normal direction. Before proceeding to the next trial, observers were instructed to verify that the shape of the surface they had implicitly indicated through the collective orientations of all of the probes appeared to match the shape of the underlying textured surface.

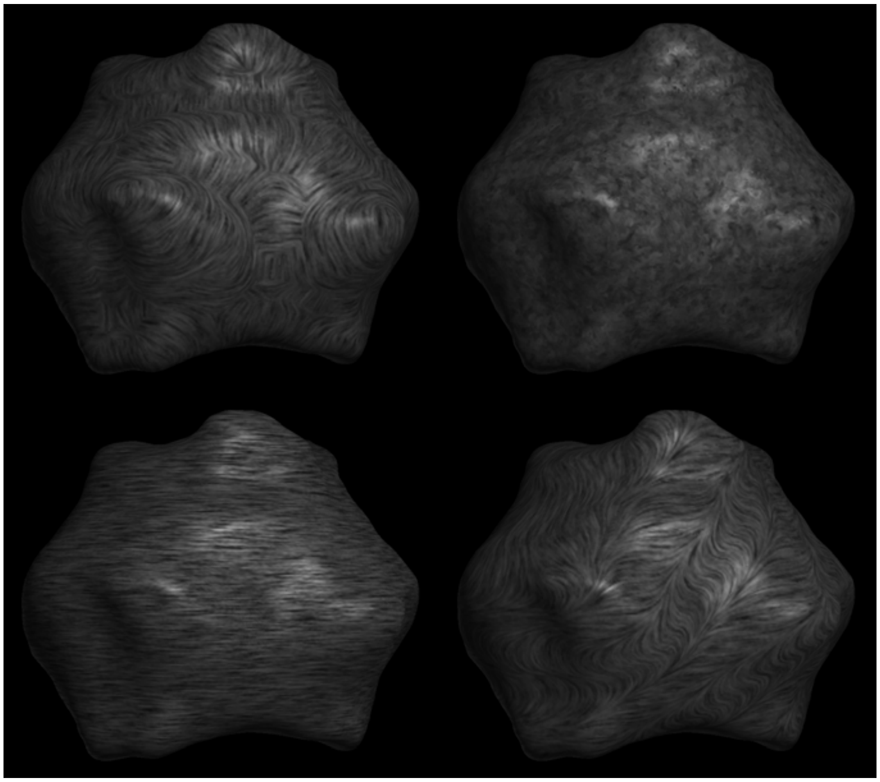

Figure 10. Examples of the 3D textured surfaces. Clockwise from top left: Principal, Isotropic, Swirly, and Uniform direction. Note that informal assessment of the potential impact of texture type on shape judgments is complicated in these images by the prominence of shape-from-contour cues, which tend to dominate when other information about shape is less readily accessible.

We prepared 24 stimuli by texturing each of the six surfaces with four different texture orientations. The stimuli were divided into two groups such that each group contained three surfaces textured in *pdir* and *sdir* orientation and three surfaces textured in *rdir* and *udir* orientation. There was no overlap between the images in the two groups. Figures 11 and 12 show the complete set of each group of stimuli.

Each observer viewed the stimuli from only one of the groups in random order. Within a group the stimuli were organized into two sets of six distinct surfaces. Observers were required to take a 10-minute rest break before viewing the second set in order to avoid any possible learning effects. The experiment was performed under conditions of flat and stereo viewing. Most of the subjects took about two hours for the entire experiment.

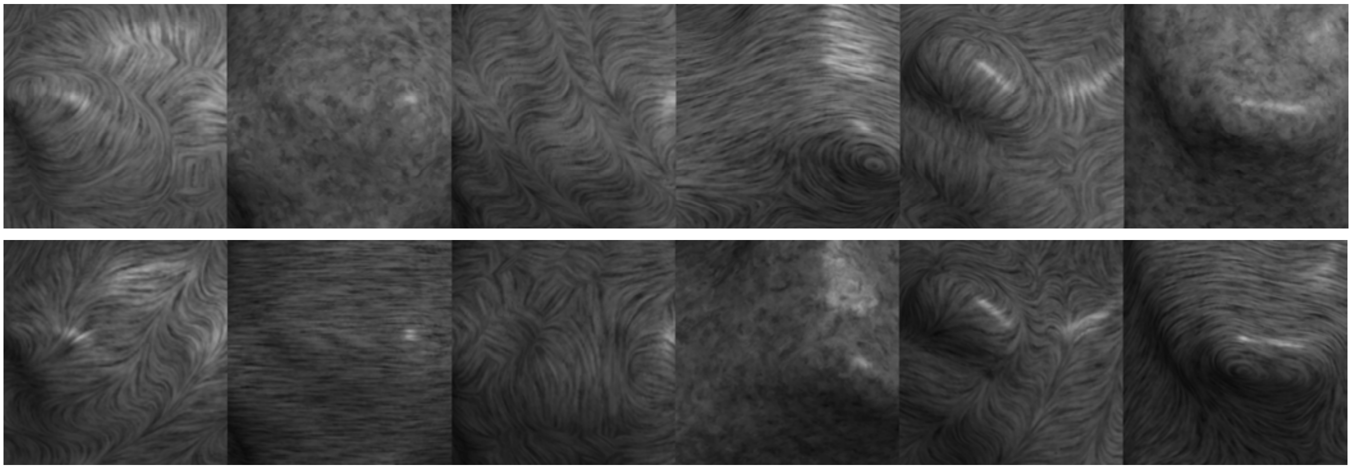

Figure 11. The set of stimuli seen by group A. The presentation order was randomly determined and was different for each subject.

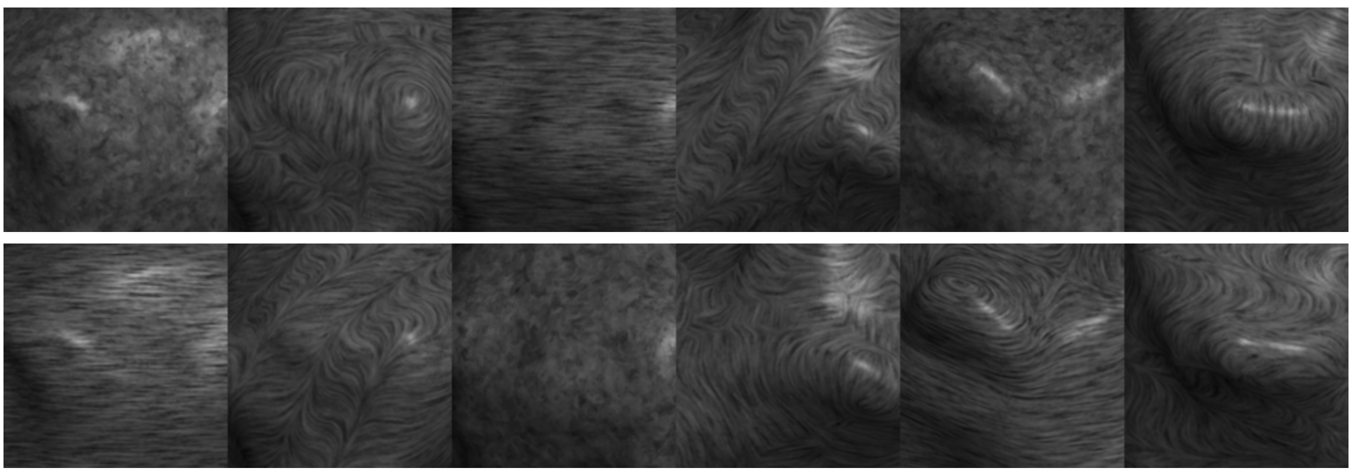

Figure 12. The set of stimuli seen by group B. The presentation order was randomly determined and was different for each subject.

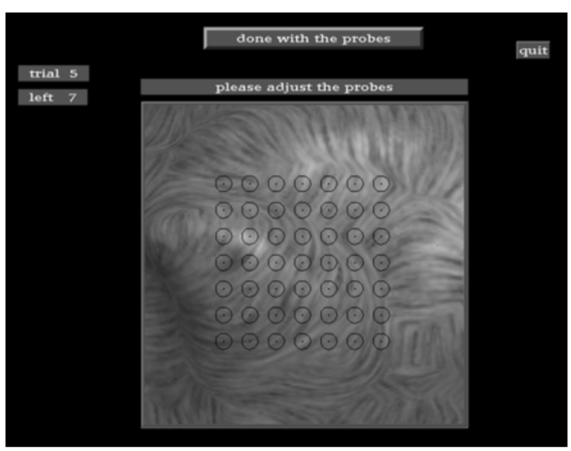

Figure 13. The graphical user interface with all probes displayed in their starting positions at the beginning of the 5th trial.

3.4 OBSERVERS

Five subjects participated in the experiment. All subjects were kept fairly naïve to the purposes of this experiment, although some of the subjects were aware of the authors' previous work with principal direction texture. We informed the subjects that we were conducting experiments to evaluate people's ability to accurately perceive 3D shape in images but we specifically did not mention anything about texture. This was done with the intent to keep the subjects as free as possible of any potential biases and to avoid leading them into certain behaviors that they might not otherwise have considered, such as lining up the direction of probe base elongation with the direction of the texture pattern, etc. Subjects were shown a single "training" image of correctly positioned probes for a surface not included in the test data and rendered without texture (Figure 14). Note that several of the probes appear to point straight out of the screen since the surface at these points is fronto-parallel.

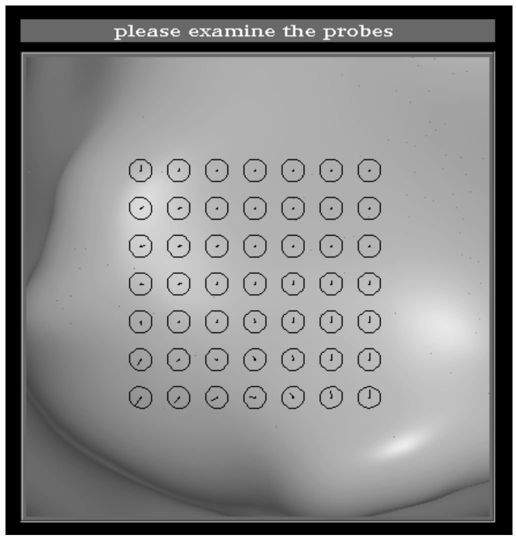

Figure 14. Training image, showing ground truth answers (correct probe orientations) at points across an untextured surface.

3.5 RESULTS

Initially we tried to measure the accuracy of observers' estimates of orientation in terms of the deviation in slant and tilt from the correct answers, where slant is the angle of rotation out of the fronto-parallel plane, and tilt is the angle of rotation about the viewing direction. Deviations in slant provided a fairly good indication of error but it was difficult to interpret the magnitude of the error due to incorrect estimates of tilt. In many cases the angular deviation in tilt was degenerate, because the estimated normal projected to a single point. However we could not simply exclude these samples from our error calculations, because their occurrence was not uniform but tended to predominate in "bad texture" conditions. In cases where the cues to shape were inadequate, subjects reverted to the default assumption that the surface lay in the plane of the image and left the probes untouched. Furthermore, even in the cases where the tilt angle was not degenerate, we felt that tilt calculations were not reliable since displacement of one or two pixels could register huge estimated errors in the tilt component.

Therefore, we used as an error metric the angle in $\Re^3$ between the estimated normal direction specified by the probe and the true surface normal direction at the probe center. Figure 15 shows the mean angular error and standard deviations computed over the 49 probe estimates under conditions of binocular flat viewing. The results are grouped by texture type, and then by test subject. Figure 16 shows the results under conditions of stereo viewing.

Under flat viewing, performance was found to be better in the cases of the isotropic pattern and the anisotropic pattern that followed the first principal direction. Under stereo viewing, accuracy increased for all texture types, but was still greater for the isotropic and principal direction patterns. Our results are consistent with a hypothesis that texture pattern anisotropy impedes surface shape perception when the direction of the anisotropy does not locally follow the direction of greatest normal curvature.

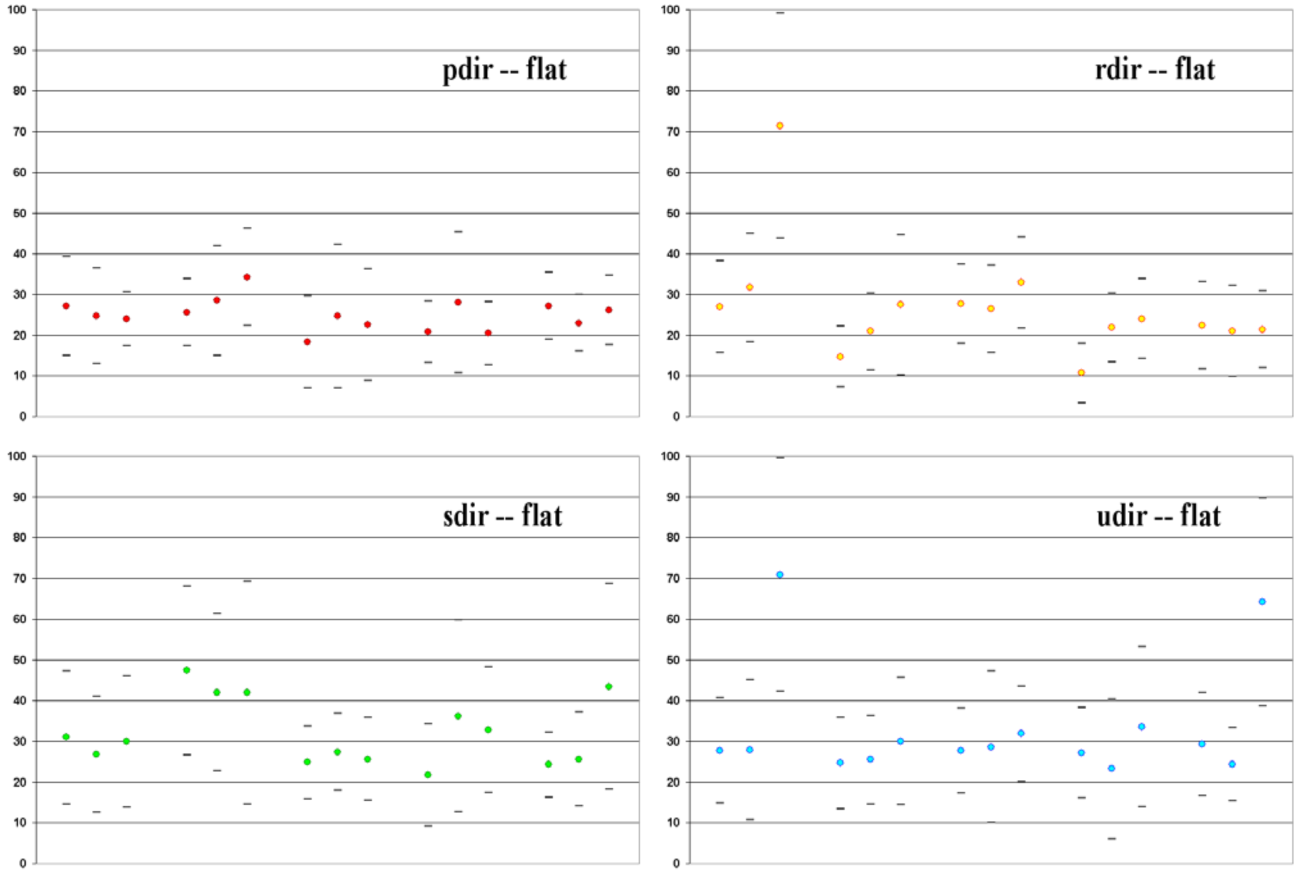

Figure 15. Individual results for the flat viewing condition. The height of each point represents mean angular error over the 49 probe locations per image. Subject number is the unspecified independent variable along the horizontal axis. Judgments from a single subject for different surfaces rendered with the same texture type are grouped by proximity along this direction. The textures are (clockwise from the top left): principal direction (*pdir*), isotropic (*rdir*), uniform (*udir*), and swirly (*sdir*).

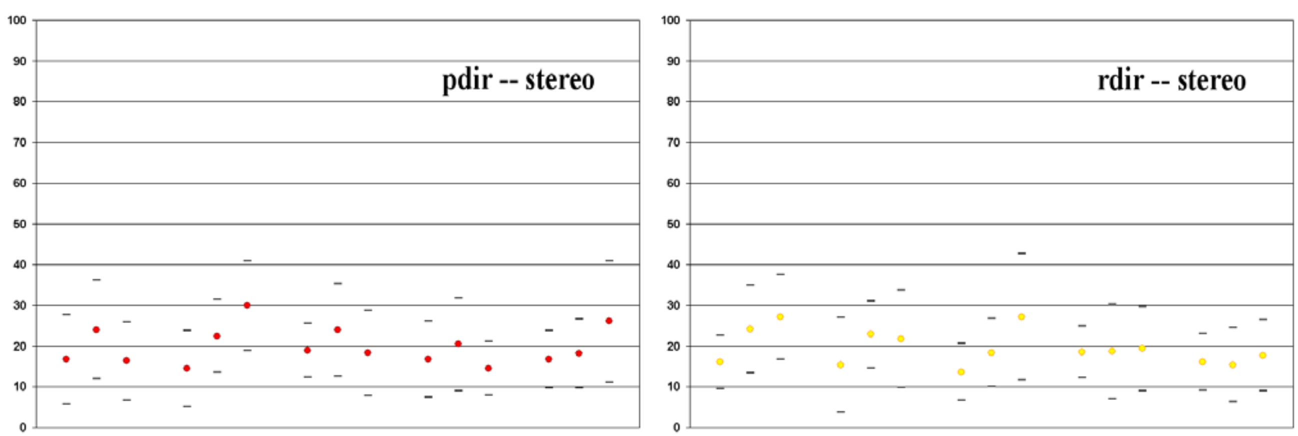

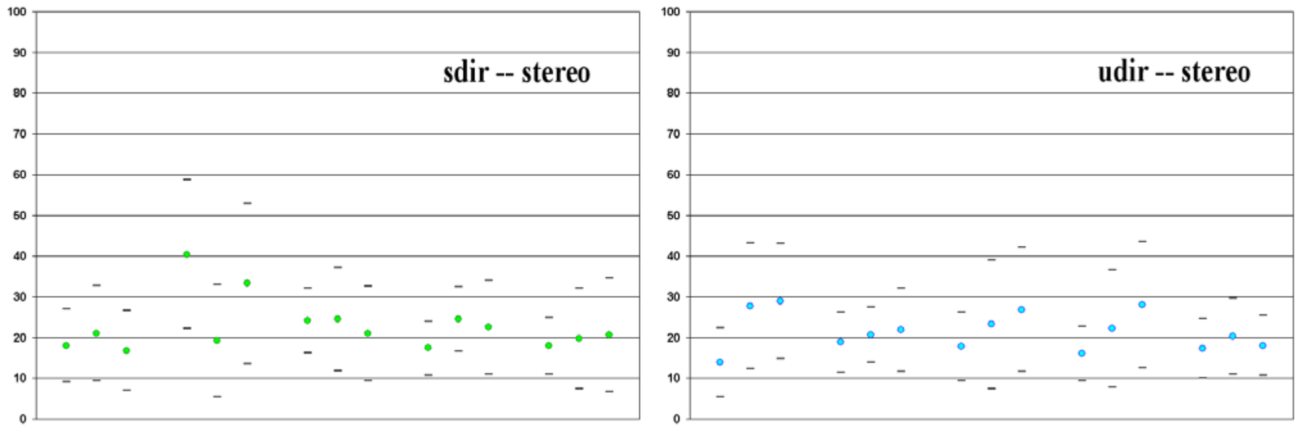

Figure 16. Individual results for the stereo viewing condition. Each point represents mean angular error over the 49 probe locations per image. Clockwise from top left: principal direction (*pdir*), isotropic (*rdir*), uniform (*udir*), swirly (*sdir*).

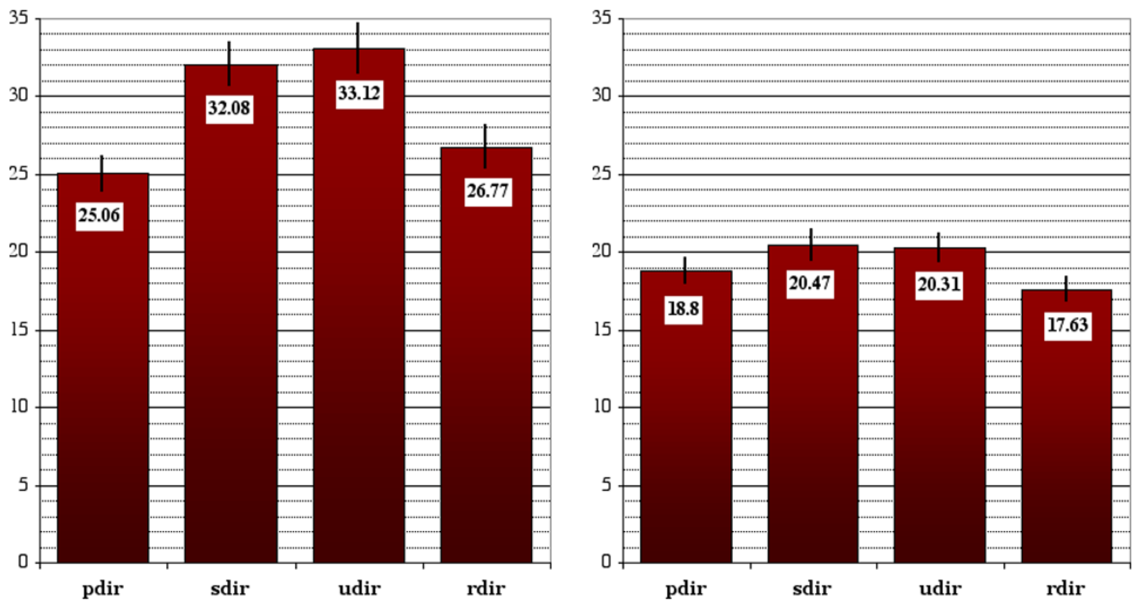

Figure 17. Pooled results (mean angular error with 95% confidence interval) for all subjects, all surfaces, by texture type. Left: flat presentation; Right stereo.

3.6 DISCUSSION

Overall, subjects seem to do better when the texture pattern is isotropic or follows the principal direction. It appears, from inspection of the individual results, that when stimuli are viewed as flat images subjects may be less prone to making catastrophic errors if the surfaces are rendered with the *pdir* texture. Closer inspection of the pattern of errors seems to suggest the presence of two different types of errors: coherent errors due to perceived depth inversion, and incoherent errors, as shown in Figure 18. Also, errors appear to accumulate in the principal direction texture around discontinuities in the pattern where the first and second principal directions switch places. We had anticipated the possibility of an advantage in using an anisotropic texture in which the direction of the anisotropy followed lines of curvature over the surface, but this interpretation is not strongly supported by the experimental results. The results suggest that the *rdir* texture did not impede shape perception more than the *pdir* texture. However some subjects were clearly misled in some places by the anisotropic patterns that followed directions different from the principal direction, suggesting that one must be careful how an anisotropic texture pattern is applied over the object if one must use an anisotropic pattern.

It is possible that our decision to place the probes at evenly spaced intervals over a rectangular grid interfered with observers' ability to perceive all of the probes as lying in the surface at the same time, due to violation of the generic viewpoint assumption. If the probes did all lie in a smooth surface that varied in depth, and still appeared to be evenly spaced in a single view, then any tiny translation of the viewing position would have to break the symmetry of the spacing. Our visual system hence preferentially adopts the more likely interpretation that the probes are placed on a transparent flat plane in front of the underlying curved surface. Our subjects did not report an inability to see the probes as lying in the surface on an individual basis, but certain of the individual responses appeared to indicate that the probes were not always consistently visualized as a coherent unit across each image.

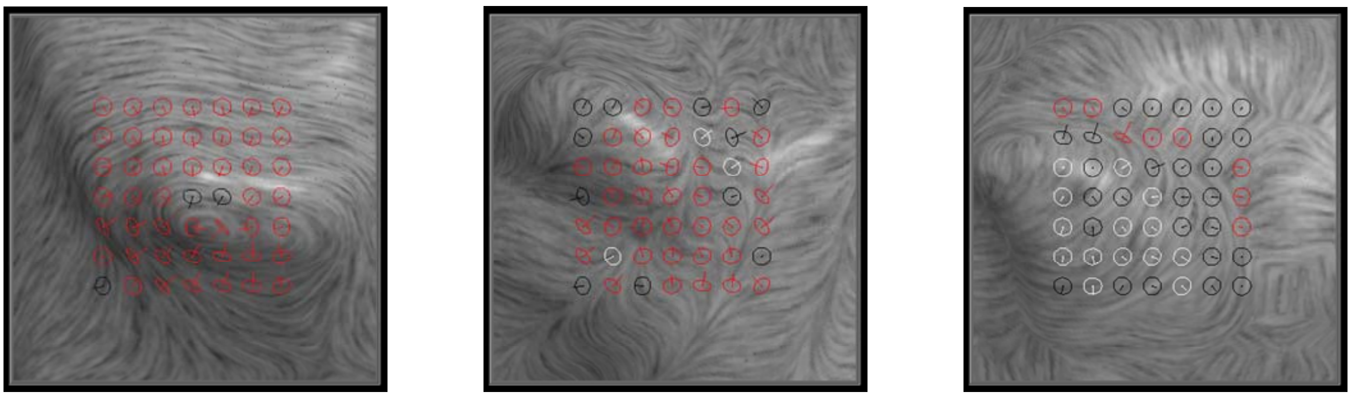

Figure 18. Some detailed individual results: Left: coherent errors due to depth inversion; Middle: incoherent errors apparently due to shape misperception; Right: errors tend to pile up at texture flow discontinuities, where the first and second principal directions switch places.

3.7 FOLLOW-UP STUDY

In a follow-up study, we repeated the experiment using displacement textures instead of luminance textures, and found the same pattern of results. Examples of stimuli for this study are shown in Figure 19.

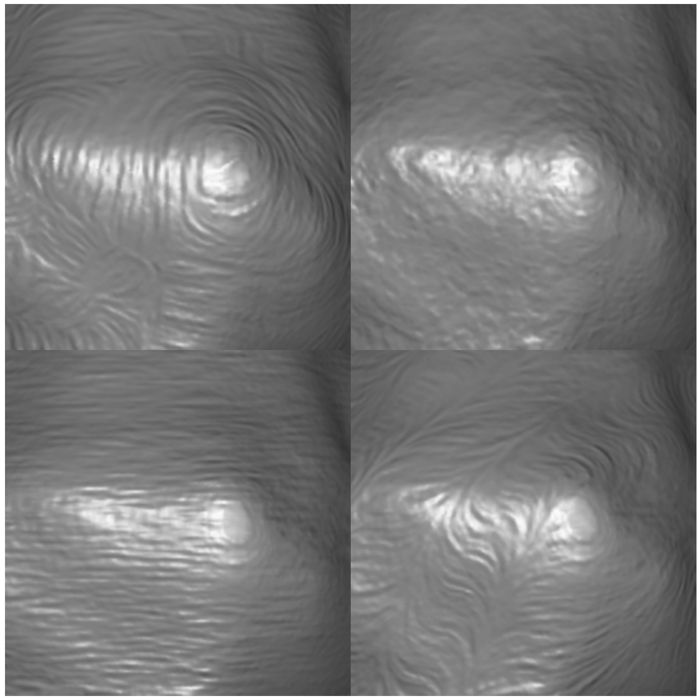

Figure 19. The same stimuli with the same textures, this time rendered as shaded relief rather than as luminance patterns.

4. EXPERIMENT 2

4.1 OVERVIEW

The results of experiment 1[11 Interrante *et al.*] using a restricted class of unidirectional texture patterns [10 Interrante] imply that: 1) accurate shape perception is most severely impeded by texture anisotropy when the flow of the texture pattern turns in the surface; 2) shape perception accuracy is not significantly different for a unidirectional texture pattern that follows the first principal direction when compared to an isotropic texture pattern of similar spatial frequency.

Two important questions were raised from the first experiment: first, if there is little ecological justification for a texture pattern to be oriented in the principal directions across a doubly curved surface, then why does shape perception seem most accurate in the principal direction orientation condition? It may be that observers are biased to interpret surface markings as being aligned with the principal directions [21 Mamassian *et al.*, 26 Stevens]. However, it is also possible that textures oriented in the principal direction intrinsically carry more shape information since they trace out lines of maximum curvature over the surface – a 3D analogy of the effect found in 2D by Biederman [3]. (From a generic viewpoint, the contours traced by a principal direction texture have the greatest potential to reveal the surface curvature to a maximum extent. For the same view, the contours traced out by a texture oriented in any other direction will be intrinsically more flat, and this may represent a loss of shape information that is not recoverable.)

Second, with arbitrary curved surfaces there are two orthogonal directions in which the normal curvature generically assumes a non-zero extrema. These directions can be reliably classified into two types known as the first principal direction and the second principal direction. It is not always clear which of these two directions should be followed by a singly-oriented directional texture in order to minimize the apparent turning of the texture pattern in the surface. Unidirectional patterns seem to exhibit artifacts at the umbilical points and parabolic lines where the first and second principal directions ‘switch places’. This raises the question whether shape perception is facilitated more effectively by an orthogonally bi-directional principal direction oriented pattern (one that has 90-degree rotational symmetry) rather than a unidirectional pattern.

In order to answer these questions we undertook a new experiment, using a shape-difference discriminability task and a new, more flexible shape-following texture synthesis method, developed in our lab, for the rendering of the stimuli. The algorithm synthesizes a high-resolution texture pattern, derived from a 2D sample, over an arbitrary doubly curved surface. The orientation of the generated texture is constrained to follow a specified underlying vector field over the surface, at a per-pixel level, without evidence of seams or projective distortion artifacts. This algorithm allows us to study in a well-controlled fashion the effects on shape perception of multiple specific texture pattern characteristics, including but not limited to orientation. Figure 20 shows some illustrative results of this new texture synthesis method that let us compare informally the shape representation efficacy of the first and second principal direction texture orientation schemes with the standard uniform direction approach. The details of our texture synthesis method are described in [8 Gorla *et al.*].

In this experiment, we attempt to use this new texture synthesis method to assess the shape information capacity of two different types of directional texture patterns (unidirectional and bi-directional) under three different orientation conditions (following the first principal direction, following a constant uniform direction, or swirling sinusoidally in the surface). In a four alternative forced choice task, we asked participants to identify the quadrant in which two B-spline surfaces differed in their shapes. The two surfaces were illuminated from different random directions and displayed simultaneously and persistently.

The two objectives of this experiment were: 1) to evaluate the validity of the hypothesis that shape perception accuracy declines if the texture does not follow the principal direction since there is less shape information; 2) to assess the advantages of using a bi-directional texture pattern rather than a unidirectional pattern in order to finesse the problem of choosing the appropriate principal direction alignment at each point. We used a four alternative forced choice surface shape discrimination task primarily for the first objective, hypothesizing that if there is less information available in the case of non-principal direction oriented patterns, then we should find a higher threshold for the detection of changes in surface shape.

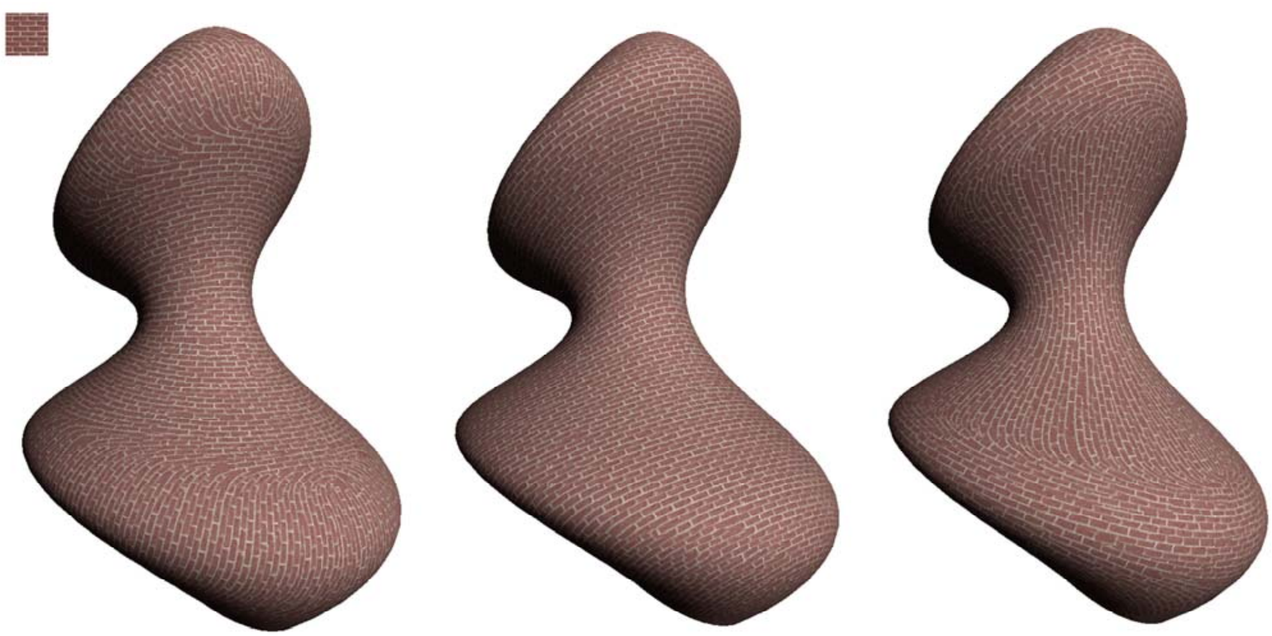

Figure 20. A brick texture pattern synthesized over a doubly curved surface according to three different texture orientation conditions. Left: rows of bricks follow the first principal direction; Center: rows of bricks are aligned in a globally uniform direction; Right: rows of bricks follow the second principal direction.

4.2 STIMULI

We used a flat B-spline surface defined by a 16 x 16 grid of control points initialized to lie at uniform intervals in x and y across the z=0 plane. We defined an initial reference surface containing randomly dispersed hills and valleys using 100 repetitions of an iterative process. At each iteration, we selected a random interior control point and displaced it by a constant amount, equivalent to 1/16th of the width of the image, in either the +z or –z direction. Then, we partitioned the reference surface into 4 quadrants, noted the control points in each quadrant that defined either a hill or a valley, and then randomly selected one of these special control points in each quadrant to control the feature that would change over the course of the trials. For each selected feature we defined 8 different displacements of the shape defining control point, in the +z direction for the hills, and in the –z direction for the valleys, and then randomly selected from among these, pairs of displacements to define 7 distinct shape difference intervals, increasing in range from 1 unit of difference to 7 units. Note that the perceptibility of a *k* unit 'shape difference' was thus equally likely to be tested with any pair of images from this set that were *k* units apart. Figure 21 shows the 8 different displacements used in quadrant 1. We were careful to compute the shading of each surface using a different random direction of illumination, selected from a solid angle of pre-determined valid illumination directions, in order to encourage our participants to attend to the 3D shape information separately conveyed in each image, and to prevent the shape difference discrimination task from degenerating into a simple 2D picture-difference discrimination task.

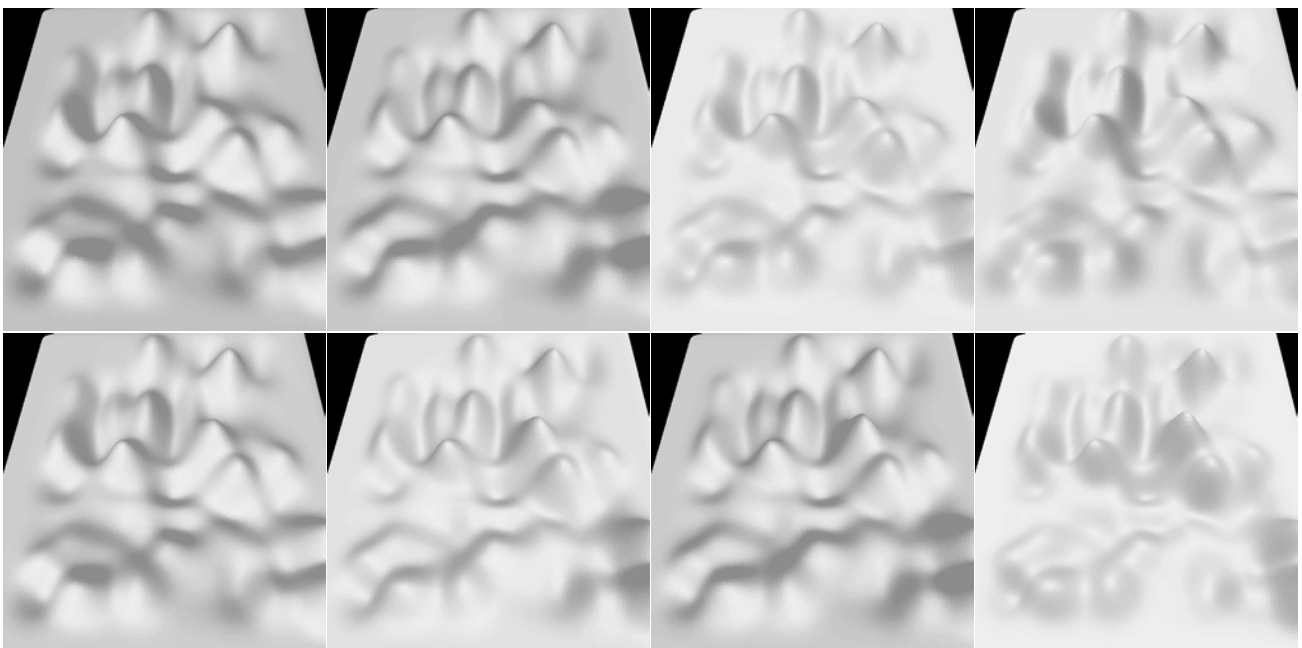

Figure 21. Shown from upper left to lower right are the eight different displacement surfaces used to represent shape changes in quadrant 1 (upper right portion of the surface).

We generated 32 sample surfaces (4 quadrants x 8 levels of displacement on the selected feature in each) and defined three different vector fields to control the orientation of the texture patterns over the surfaces. We wanted to compare three different texture orientation conditions: 1) the principal direction condition, in which the texture was constrained to follow one or both of the principal directions at every point over the surface; 2) the uniform direction condition, in which the texture was constrained to remain oriented in a constant uniform direction across the surface; 3) the swirly direction condition, in which the orientation of the texture pattern twisted and turned in a sinusoidal manner over the surface. Since we had the parametric definition of the B-spline surfaces, we were able to compute the first and second principal directions analytically at every vertex in the mesh. We obtained the uniform direction vector field by calculating the line of intersection of the tangent plane at every vertex with the plane that passed through that vertex in an orientation parallel to the x=y direction. Finally, we obtained the swirly direction vector field by rotating each of the uniform direction vectors by a coherently varying amount that was defined as a continuous function of the 3D position of the vertex in world coordinate space.

All surfaces were textured using the 'fitted texture' synthesis program. We used two base texture patterns selected from the Brodatz texture album [2]. We chose a unidirectional pattern D49: Straw Screening and a bi-directional pattern D20: French Canvas. We divided each of these patterns into eight subimages, so that the surface texture could be created from a different but similar pattern for each of the different surface displacement intervals. This allowed us to avoid unwanted pixel-by-pixel texture similarities in areas of the surface that did not undergo a shape change.

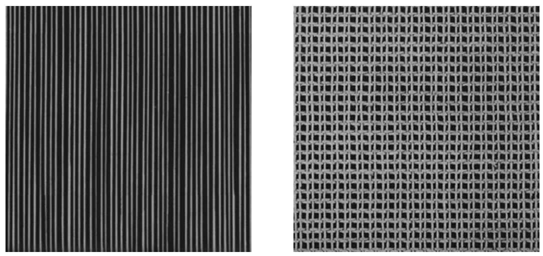

Figure 22. D49 and D20 from the Brodatz texture album.

Figure 23 shows the same series of surfaces presented in Figure 21, with the bi-directional texture pattern applied following the orthogonal principal directions. Figure 24 shows the results of using the three different texture orientation conditions (uniform, principal and swirly) over the same underlying surface.

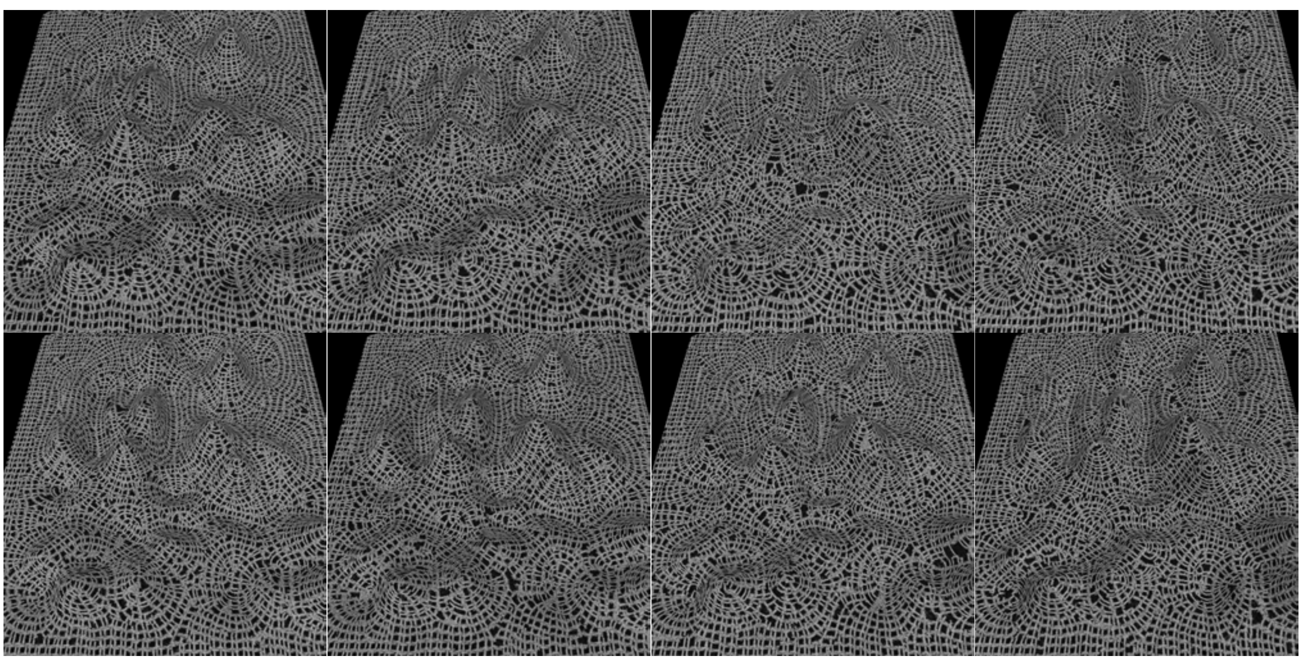

Figure 23. From upper left to lower right, the same eight different displacement surfaces, textured with principal direction oriented patterns derived from eight different subsample swatches from the same larger original texture pattern image.

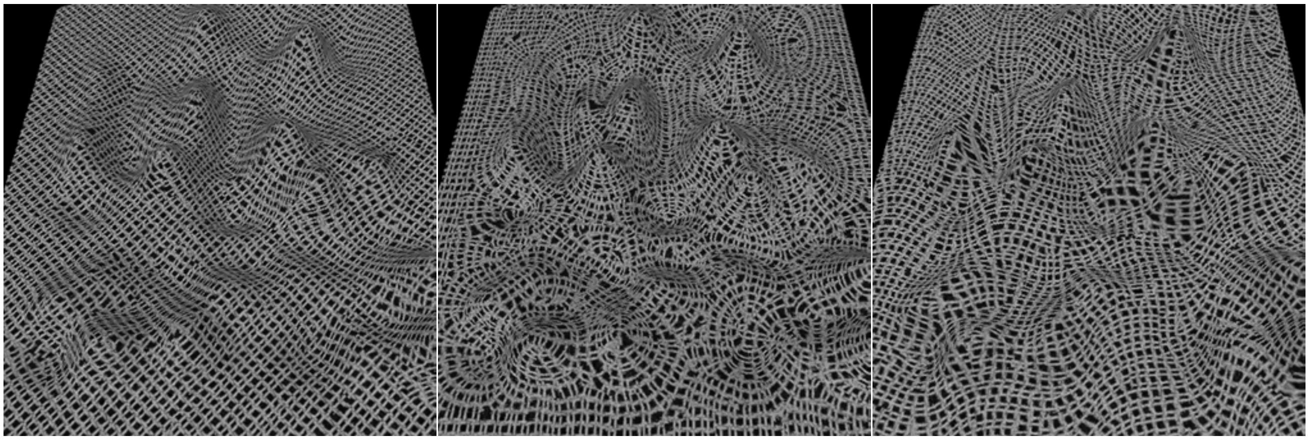

Figure 24. The three different texture orientation conditions. Left: uniform direction; Center: principal direction, Right: swirly direction.

In order to avoid the possibility of any systematic biases due to the choice of viewing angle, we used two different viewing conditions for all surfaces: frontal and tilted. Examples of the frontal view and the unidirectional texture can be seen in the Figures in the following section.

4.3 EXPERIMENTAL SETUP AND TASK

The experiment consisted of 672 trials (3 orientation conditions x 2 texture pattern conditions x 4 shape types/quadrants x 7 shape displacement amounts x 2 viewing conditions x 2 repeated measures). Observers were shown pairs of images, side-by-side, and asked to specify in which quadrant the surface shapes appeared to be different. Viewing time was unrestricted. In analyzing the data, we found that subjects took, on average, between 14-25 seconds to make their decisions, spending a total of between 2.7 – 4.7 hours overall, including breaks. Figures 25 and 26 show the user interface for the experiment with representative examples of the remaining conditions. The users were able to display or hide the boundaries between the quadrants over the image (the boundaries were defined in 3D).

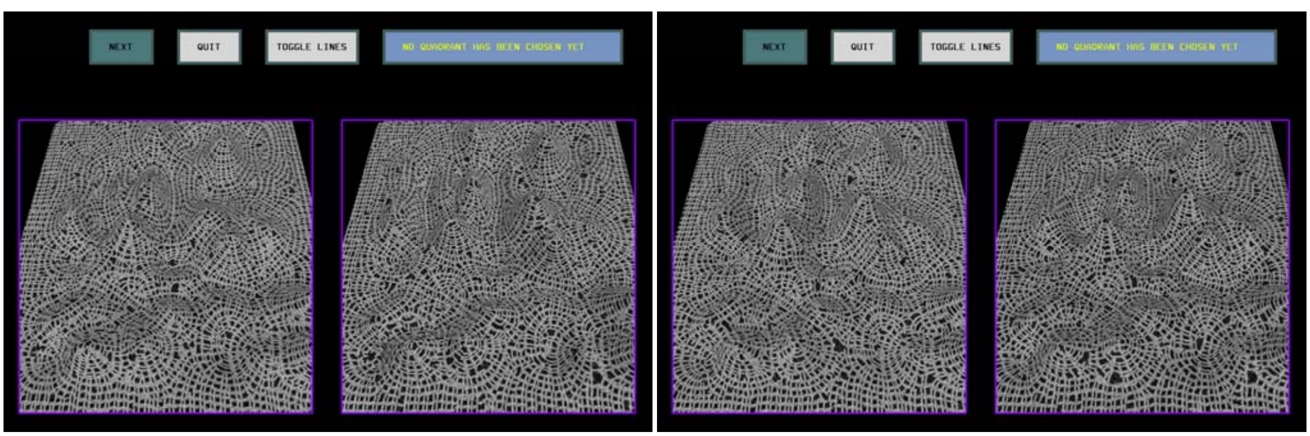

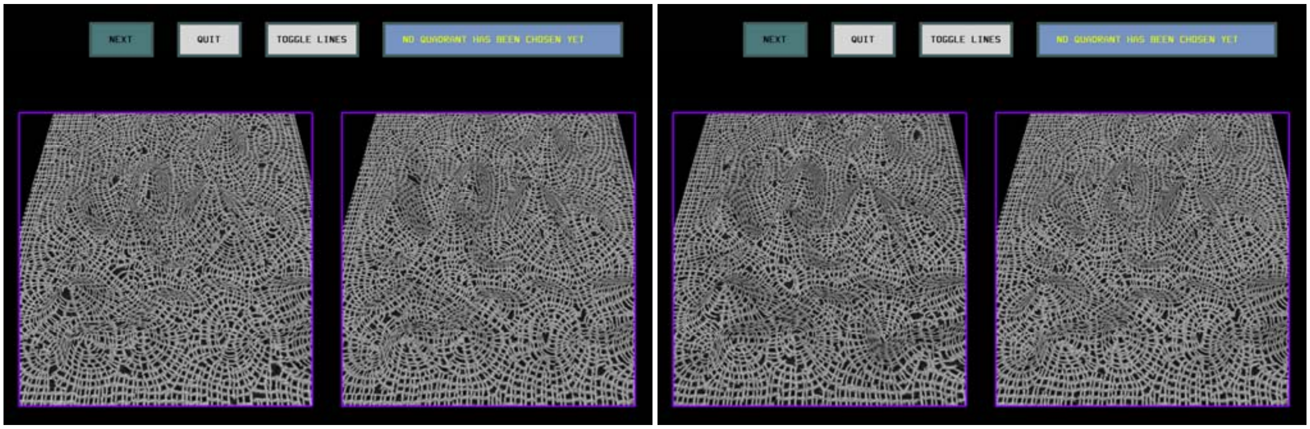

Figure 25. Screen shots from the experiment, showing the shape difference in each of the four quadrants. The bi-directional pattern and the principal direction orientation condition were used in all of the cases shown here. Clockwise from the upper left, the surface shape differences appear in the following quadrants of each image respectively: upper right, upper left, lower right and lower left quadrant.

We collected data from three subjects. One of the principal investigators participated as a subject and was intimately familiar with both the task and the experimental objectives but was not involved in the preparation of the stimuli. The other two subjects were professionals from outside of computer science, who were new to the task but informally aware of our basic research objectives.

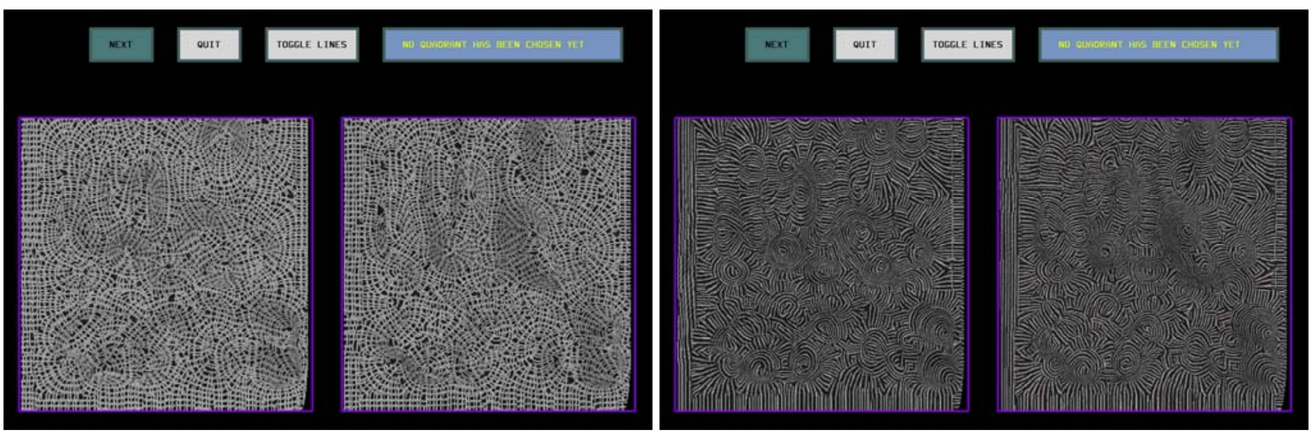

Figure 26. Examples of the two remaining conditions: the frontal view, and, on the right, the unidirectional texture pattern in the first principal direction orientation condition.

4.4 RESULTS

After all subjects had gained sufficient training in the task, their accuracy increased fairly consistently with increasing magnitude of shape disparity, but the characteristics of this increase differed under the different texture orientation conditions. Subjects were able to perceive smaller shape differences most reliably when the surfaces were textured with a pattern whose orientation followed one of the principal directions. Their accuracy for small shape differences diminished when the surface textures either gradually swirled in the surface or followed a constant uniform direction in the tangent plane regardless of the shape characteristics. These findings appear to support our hypothesis that anisotropic textures aligned with the first principal direction may facilitate shape perception for a generic view. First principal direction patterns appear to convey to the observer more reliable information about the extent of the surface curvature than texture pattern oriented in any other direction.

Figure 27 summarizes the overall results of our experiment. Each data point in these images represents the percentage of correct quadrant selections that each user made over the 32 different trials corresponding to each shape difference level, for each texture orientation condition. (Two repeated measures were taken for each pair. Our graphs show the combined data across the two viewing conditions, the two texture pattern conditions, and the four quadrant/shape types.) We did not consider each of the texture pattern conditions separately because the results turned out to be similar in both cases. The pattern of responses seems similar across the three subjects, although the level of performance seems higher, overall, for the expert subject. Accuracy rates rise as the disparity of the compared shapes increases, with a faster, earlier increase in the case of the principal direction textures, and a more linear but slower increase in the case of the swirly and the uniform direction textures. Based on consistent reports from all three subjects that "the task got a lot easier once you figured out what to look for", we decided to consider the first and second half of the data separately. Figures 28-29 show the results tabulated separately for the first 336 responses and the second 336 responses per subject.

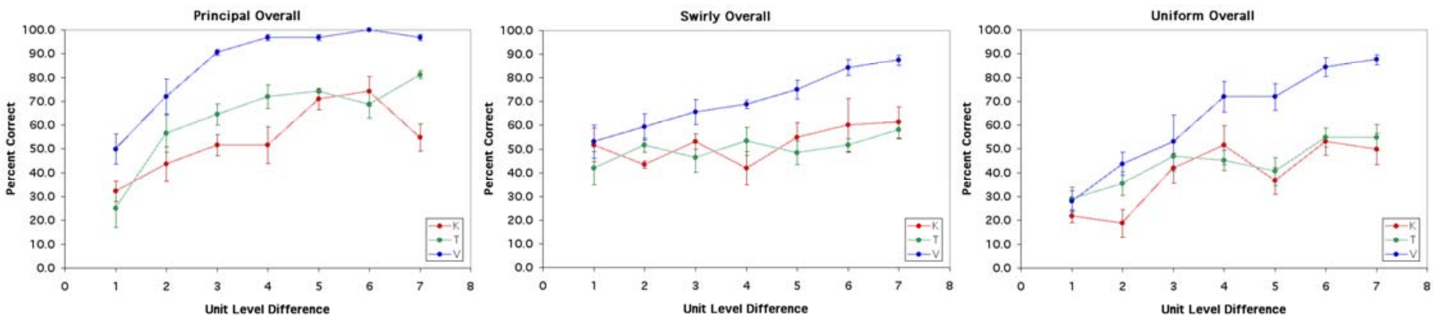

Figure 27. Overall results from the experiment, per texture orientation condition. Each line represents a different subject.

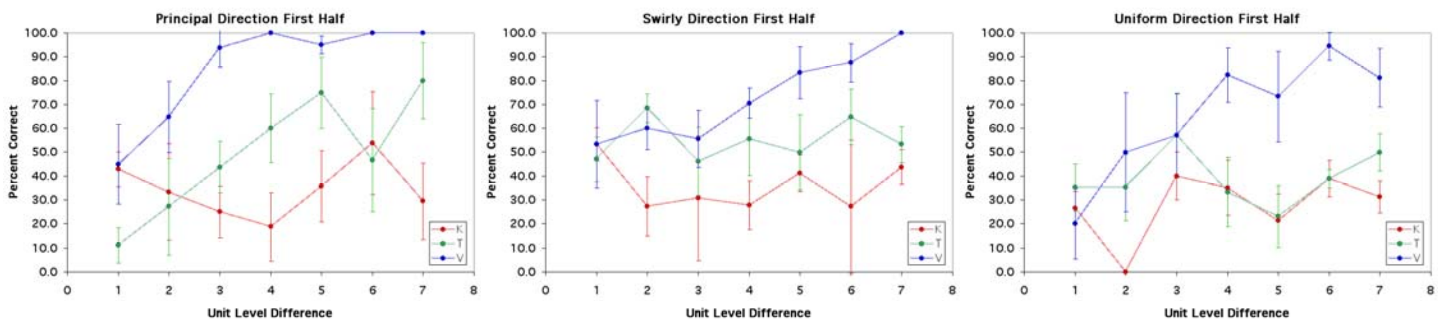

Figure 28. Results from the 'first half' of the trials only, per texture orientation condition.

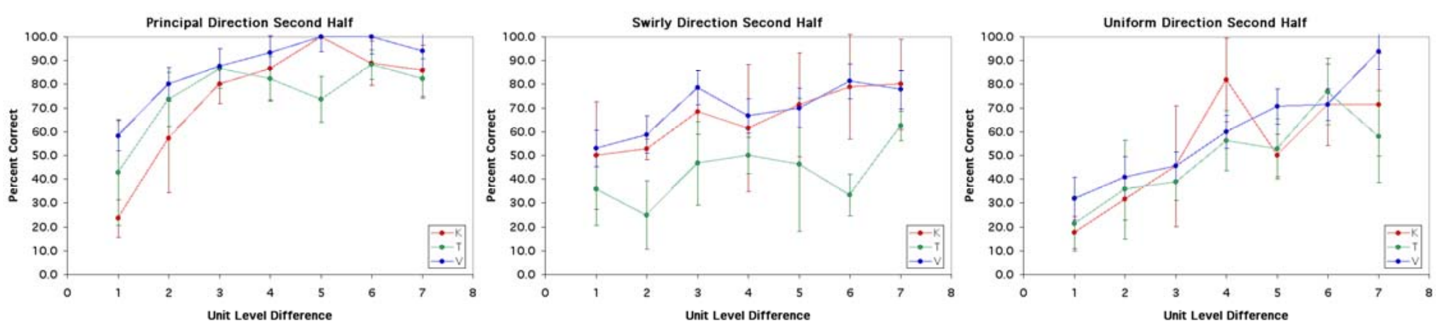

Figure 29. Results from the 'second half' only, per texture orientation condition.

It is clear from the data in Figures 28 and 29 that the performance of the three observers converged well in the second half, but poorly in the first. We suspect that in the beginning of the experiment the observers were unclear about the assigned task. Since only four distinct shape features actually changed, it is likely that all of the observers eventually realized how to differentiate between the apparent image differences that were due to the vagaries of the texture synthesis process and the differences caused by subtle changes in the shapes of the underlying surfaces. After this point, the observers' task performance may have been more closely tied to the actual availability of shape information and less sensitive to distraction by irrelevant texture variations.

In Figure 30, we combine the performance results across the displacement intervals to get a rough comparative overview of the overall task performance under the three different texture orientation conditions. We note that the expert observer, *V*, seems to display fairly consistent performance across the two halves of the experiment, while naïve participant *K* seems to exhibit a tremendous improvement in performance from levels just above chance in the first half to levels equivalent to those achieved by *V* in the second half, across all three texture orientation conditions. Participant *T* seems to show a dramatic improvement in performance only for the principal direction textures, simultaneously accompanied by an apparent *decrease* in performance in the case of the swirly textures.

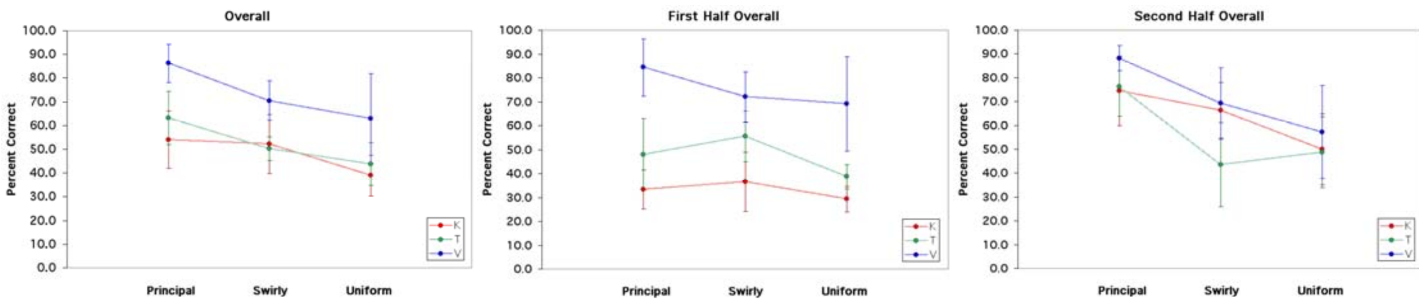

Figure 30. Combined results (averaged over all displacement ranges), per texture orientation condition.

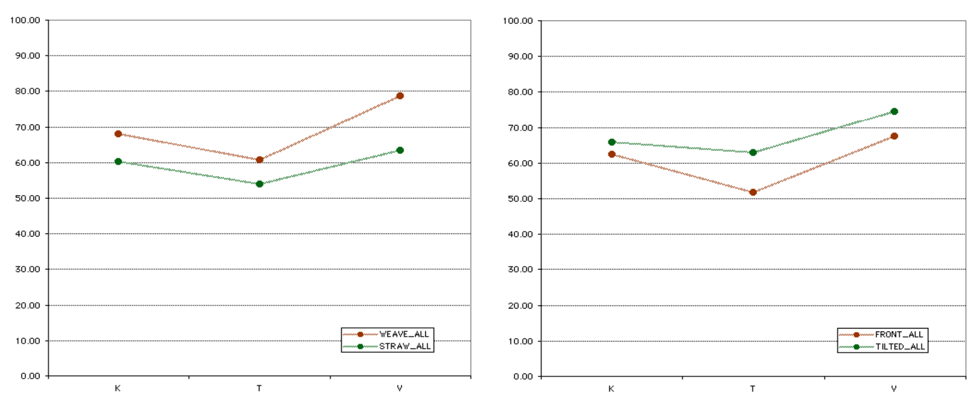

Figure 31. A comparison of performance under the two texture type conditions (left) and viewing conditions (right).

Figure 31 summarizes the overall differences, or lack thereof, in the cases of the two texture type conditions and the two viewing conditions. Performance appears to be slightly better, on average, for the tilted surface position compared to the forward-facing position, probably because it is easier to gauge the heights of the peaks under the tilted condition. Performance also appears slightly better for the bi-directional texture than the unidirectional texture, but the difference is statistically insignificant.

4.5 DISCUSSION

The results from this experiment seem to lend support to the hypothesis that principal direction oriented texture patterns have a potential advantage over non-principal direction oriented patterns in facilitating shape perception because of the fact that they provide more explicit evidence of the surface curvature than do directional patterns oriented in any other way. Disappointingly, but in hindsight perhaps predictably, the experiment did not provide us with much insight into the second of our objective queries – to assess the relative potential advantage for shape representation offered by a texture pattern that explicitly followed both the first and second principal directions simultaneously rather than following only one of these directions.

One question that is raised by this experiment is, why we seemed to find, in general, that task performance was worse in the uniform direction condition than in the swirly condition? Theoretically we would have expected that the insurface undulations of the swirly texture would have made things worse for shape understanding. The answer, we believe, lies in recognizing that task performance ultimately depended not on the accuracy of the shape understanding but only upon the discriminability of the presence of shape changes. All of the shape displacements were due to the translation of selected control points in a direction orthogonal to the original base plane. Since the elongated elements in the uniformly oriented textures were defined to lie in a constant direction parallel to the x=y plane, stretching the surface in the z direction had remarkably little effect on the appearance of the texture pattern. Surfaces that were quite different in shape had uniform direction vector fields that were nearly indistinguishable when viewed from above. As we noted earlier, this is an intrinsic problem with uniform direction textures. They do not depend much on the surface geometry, and hence there will inevitably be some aspects of a surface's shape that from some viewpoint such textures are not well equipped to represent.

4.6 LIMITATIONS

In an attempt to minimize the possibility of inadvertently biasing the experiment through unconscious interference, we were careful to ensure that the entire surface definition process was done in a fully automated way without any manual intervention. Unfortunately we neglected to foresee the possibility that some shape differences might be partially hidden from view of the observer due to occlusion by other parts of the surface. In many of the tilted views of our surfaces the bottom-most tip of the valley feature in quadrant 0 is not fully visible to the observer. We do not believe that our overall findings were biased by this occurrence. However we feel that for completeness this observation should be noted.

It is possible that our decision to vary the illumination direction randomly may have had some unintended and undesirable consequences. It is well known that the accuracy of shape perception judgments vary under different illumination conditions. In particular, it has been shown that shape perception is facilitated to a greater extent when the incident light strikes the surface at a grazing angle rather than when the light hits the surface head-on [13 Johnston]. By varying the direction of illumination randomly for all surfaces, it is possible that we inadvertently made the shape discrimination task more or less difficult in some cases than in others, as a consequence of the illumination direction. Because of the high number of trials (225 judgments per texture condition), and the subtlety of the lighting effects on top of the texture, we believe that it is unlikely that the overall final findings have been significantly biased because of this issue. However, it is worth considering whether a different strategy might be used, to avoid this problem and potentially obtain more reliable results.

Finally, a principal rule in good experimental design is that only one independent variable should change at a time. This means that if we are going to compare the discriminability of one unit of shape difference between two surfaces under three different texture conditions, we would like to be sure that the only difference in these three cases is the orientation of texture over the surface. The random selection of lighting direction and random selection of difference interval endpoints must be the same for all three texture conditions in that particular one unit difference case. Unfortunately after the experiment was completed, we discovered that the random selection feature for both lighting and interval endpoint choice was left on throughout all trials, resulting in the situation that the shape difference discrimination task was *not* performed under strictly identical conditions across all texture types. We believe that the conditions were in all respects equivalent, and we do not believe that any systematic bias was introduced through these random variations, especially because of the large number of trials. However it will be important for us to re-validate our findings in the immediate future with a follow-up experiment in which these extraneous random variations are more carefully controlled for.

While it is logical to assume that observers' ability to perceive surface shape accurately is poor when they are unable to identify accurately the quadrant of an image in which two presented surfaces differ in shape, the inverse assumption is not well founded. One cannot infer that in instances where observers are able to identify accurately the quadrant in which two presented surfaces differ in shape, that this is because they are in fact able to perceive the two shapes accurately. It is quite easy to imagine a scenario in which a subject has an invalid interpretation of the shapes of each of two surfaces, but can still discern that the two surfaces have different shapes. Hence, the conclusions drawn from this study can be used only to inform our understanding of the *potential* of textures of various orientations to carry shape information. It is not surprising that the principal direction textures would exhibit more prominent variation in response to shape changes than the other two types of textures. This is because the principal direction textures are defined according to surface shape properties while the relationship between the other two textures and the surface shape is more indirect. Further investigations are needed to probe the extent to which shape understanding is facilitated under different texture type conditions.

5. EXPERIMENT 3

5.1 OVERVIEW

From the first and second experiments, we believe that the orientation characteristics of an anisotropic texture pattern matter [11, 12]. We now address the remaining important questions about how best to determine the characteristics of a subtle and aesthetic texture pattern that can be used to facilitate veridical shape perception without introducing unwanted visual noise. If we want to use a principal direction oriented pattern, what *kind* of principal direction texture is best? Does it matter? How can we characterize what helps? We also wish to investigate if a texture pattern that follows both first and second principal directions conveys the shape of the underlying surface better than a texture pattern that follows only one of the principal directions.

In experiment 3, we compared performance of a surface attitude probe adjustment task (introduced in experiment 1) to evaluate the relative effectiveness of three distinct conditions of principal direction pattern orientation and a control condition in which no texture was present. The three texture conditions were: 1) a doubly-oriented texture in which approximately evenly-spaced lines follow both of the principal directions ('2-dir'); 2) a singly-oriented line texture which follows only the first principal direction ('1-dir'); 3) a singly-oriented line integral convolution texture ('lic'), from which information about texture compression in the direction of the texture flow may be indirectly accessible.

The immediate goal of the third experiment was to determine whether some principal direction texture conditions allow observers to make more accurate surface shape judgments. In order to make the problem tractable, we chose patterns that varied in only one or two important respects. Specifically we were interested in determining whether shape perception might be facilitated in the condition of a texture that contains elongated elements that can be interpreted to follow both of the principal directions simultaneously compared with a texture in which the elongated elements are oriented solely in the first principal directions. Additionally, we were interested in probing the potential effects of other texture pattern characteristics, besides orientation. For this reason we decided to also evaluate shape perception in the case of a singly-oriented line integral convolution-like pattern ('lic').

There were 200 trials (4 texture conditions x 10 surface/probe locations x 5 repeated measures). A total of five naïve participants were asked to adjust a circular probe, randomly located on an arbitrary doubly curved surface, so that its perpendicular extension appeared to be oriented in the direction of the surface normal.

5.2 STIMULI

We first defined the texture samples that characterized the patterns that would be synthesized over the surface stimuli. Using Inklination's Pen-and-Ink Crosshatching Filter plug-in for Adobe Photoshop, we created the two-directional and one-directional patterns shown in the left and center of Figure 32 from the same uniform light gray base pattern. The three patterns have nearly equal mean luminance (2-dir=156.26, 1-dir=159.69, lic=127.04), but the line patterns have significantly different luminance histograms and standard deviations from the lic pattern (2-dir=84.37, 2-dir=97.93, lic=21.10).

Figure 32. The texture patterns used in the study. From left to right: two-directional (2-dir), one-directional (1-dir) and lic-like (lic).

Next, we defined arbitrary smoothly curving surfaces that the participants would use in making their surface shape judgments. Figure 33 shows the textures on one of these test surfaces. We used the procedures from experiment 2 to create the surfaces. Having the parametric definition of the B-spline surfaces, we were able to compute the first and second principal directions analytically at every vertex in the surface mesh to use in the texture synthesis.

In the final step in the preparation of the stimuli, we used the "fitted texture" synthesis to define actual surface texture. We synthesized each of the three test patterns over each of the five test surfaces shown in Figure 34.

5.3 EXPERIMENTAL SETUP AND TASK DESCRIPTION

After defining the surfaces, we selected the locations of the surface attitude probes. For consistency between observers we pre-determined a fixed set of ten locations, two on each surface, at which the users would make surface orientation judgments. It was essential to the integrity of the experiment that the probe locations be determined completely randomly, in order to avoid inadvertently biasing the results through an unconscious preferential selection of positions at which the shape appeared "well-behaved" or comprehensible. However, we did reject probe positions that were not visible from the predetermined viewpoint, and probe positions at which the default initial probe orientation was within 10 degrees of the true surface normal direction, in which case participants would be able to get "good" results without performing any task.

We limited the study to ten probe positions to control for fatigue-related factors and to keep each session within two hours. Stimuli were displayed in a 900x900 pixel window on a 21" Sony Trinitron Multiscan E500 monitor and freely viewed from an approximate distance of 24". Both the surface and the probe were modeled in 3D and displayed in perspective projection using an OpenGL based renderer, and shaded using a standard Phong illumination model. The viewing angle and lighting parameters were held constant over all trials. Observers could freely rotate the probe in 3D by clicking and dragging the mouse in a way that simulated the effect of pulling on the probe handle. However, observers could not get any occlusion cues while manipulating the probe.

The procedure to determine the presentation order of the 200 trials was refined through a small pilot study involving two of the principal investigators. We determined that because of the greater ambiguity of the local surface orientation in the untextured condition, and the strong incentive to carry over inferred surface orientation information gleaned from textured trials, it would be necessary to have participants fully complete the portion of the experiment involving the untextured trials before proceeding to any trials in which the surface was textured. In addition, we created a method to avoid the potential bunching up of presentations involving repeated measures at any individual probe location.

Before starting the experiment, participants were given an instruction sheet, which contained the explanation of the experiment and their task. We provided written instructions in order to minimize the chances of our inadvertently coaching different participants in different ways. Five Computer Science students participated in the experiment.

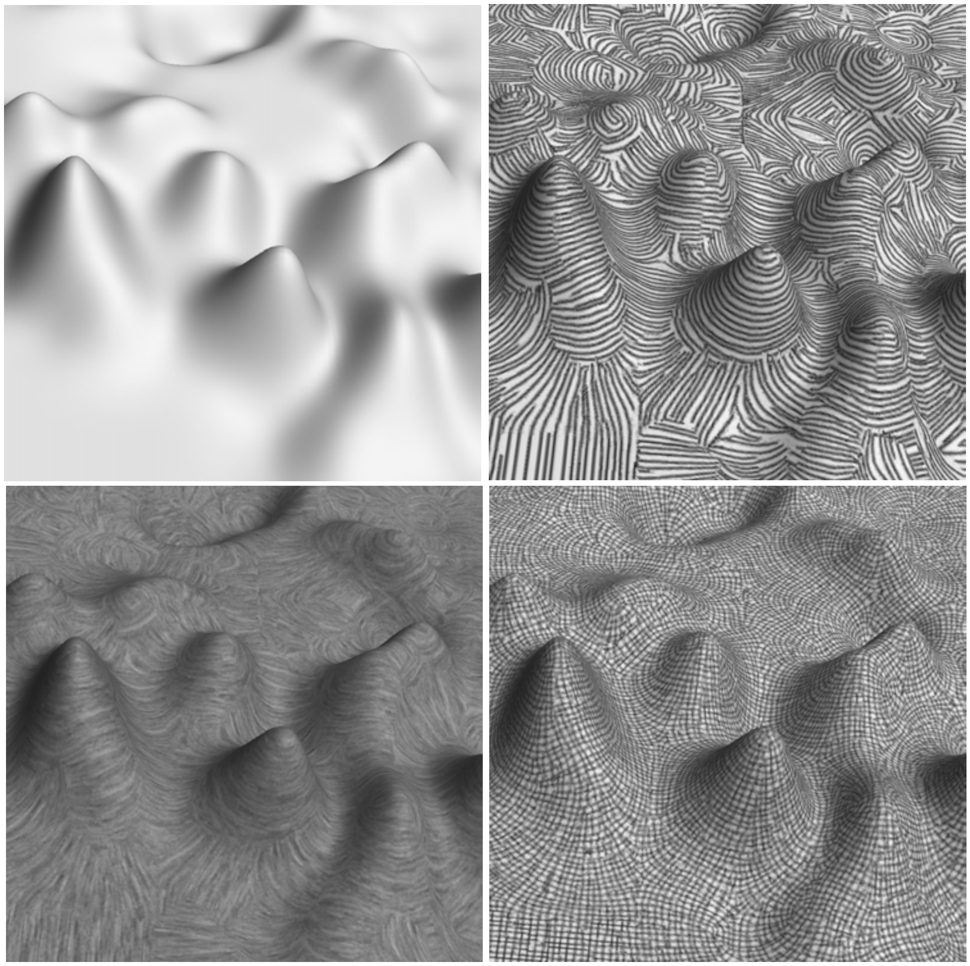

Figure 33. The four texture types used in our experiments. Upper left: No texture (control condition). Upper right: One-directional line pattern, following the first principal direction. Lower left: One-directional LIC pattern, following the first principal direction. Lower right: Two-directional line pattern, following the first and second principal directions.

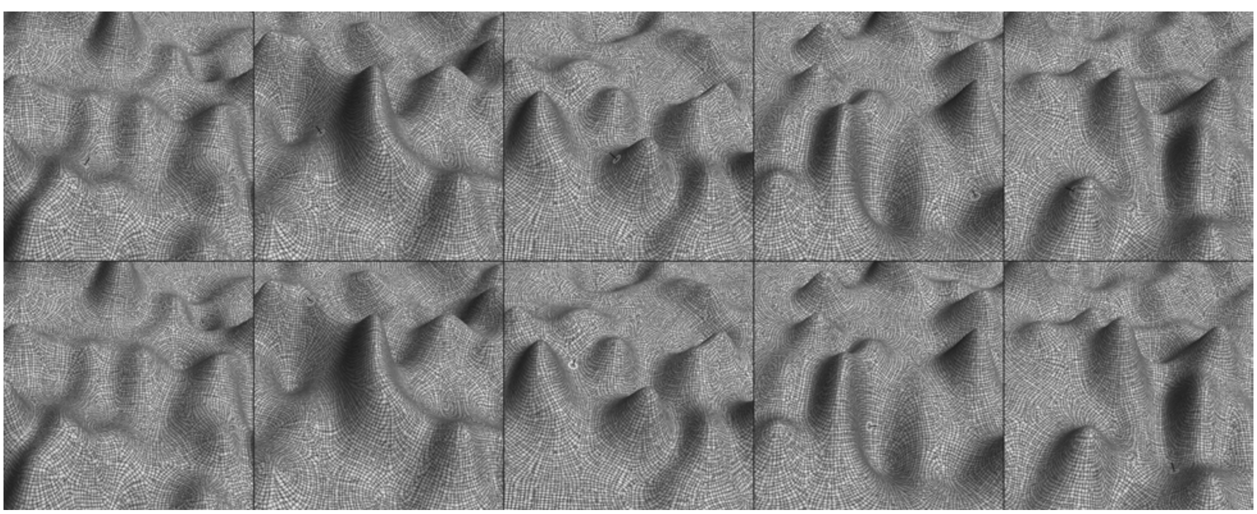

Figure 34. The five surfaces and ten probe positions used in our study. Even numbered probes 0,2,4,6,8 appear in the top row, with the odd numbered probes below them.

5.4 TRAINING PHASE

To ensure that participants had an adequate understanding of the task, each participant was required to complete a training session immediately before the experiment. We generated a sixth surface for the training, which was textured using an isotropic random noise texture that was close to the test patterns in mean luminance and spatial frequency. As in the actual study, a probe was superimposed on top of the textured surface in each trial. Only one surface was used, with 15 different probe locations individually presented. Subjects were asked to manipulate the probe until they were satisfied that the probe's perpendicular extension had the same direction as the normal to the surface at that point. After pushing the "NEXT" button, if the probe orientation selected by the user was within 10 degree of the true surface normal orientation, they would automatically proceed to the next trial. Otherwise, the probe would be color-coded based on the magnitude of the error, measured as the three dimensional angle between the true normal and the user selected probe normal. At this point, the user would be able to continue with the probe manipulation until they had determined an adequately accurate position for the probe, assisted by the color as a cue to correct their estimates dynamically. In order to prevent users from relying 100% on the color-coding without actually trying to understand the shape of the surface, we required that each subject pass three trials out of the 15 without using the color-coding cue. Figure 35 shows an example of the training data.

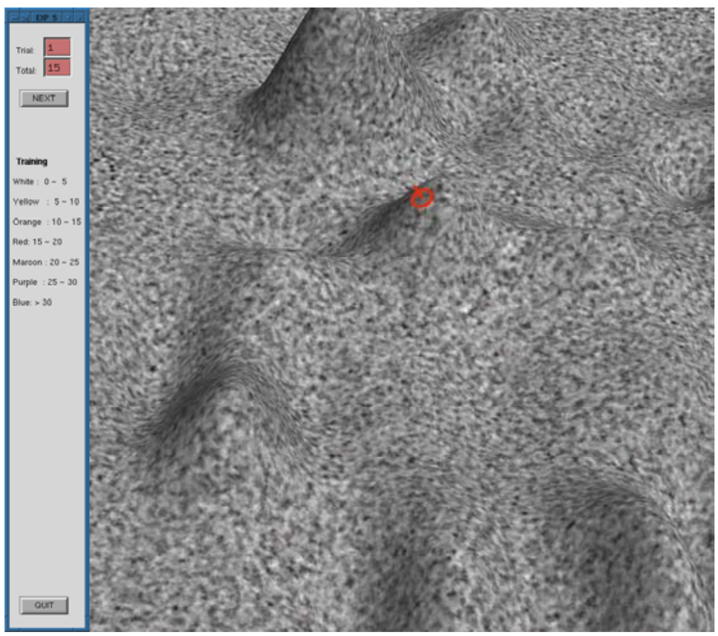

Figure 35. The training surface with one of the 15 test probes, shown in an orientation that is within 15-20 degrees of the true position.

5.5 RESULTS

Figure 36 shows an overall summary of the results that we found in this experiment. Both the mean and median angle errors, across all observers and all probe locations, had the same pattern. Performance was best in the case of the two-directional pattern, closely followed by the lic-like pattern, and then the one-directional pattern. As expected, performance was worst in the no texture condition. We used the statistical software package 'MacAnova' to perform a three-way, within-subjects mixed analysis of variance (ANOVA) to evaluate the statistical significance of the results. We found significant effects of probe location ($p=0.0000264$) and texture type ($p=0.0002843$), and a significant two-way interaction between texture type and probe location ($p<0.00000001$). We did not find a significant effect of subject id ($p = 0.18$) nor of a significant interaction between subject and texture type ($p = 0.62$). We used Tukey's HSD ("Honestly Significant Difference") method to perform post-hoc pairwise comparisons of the means of the angle errors under the different texture conditions. We found that the following differences were statistically significant at the 0.01 level: 2-dir < 1-dir, 2-dir

< None, 1-dir < None, and LIC < None. The difference between performance in the 2-dir and LIC conditions was not statistically significant at the 0.01 level, nor was the performance difference between the LIC and 1-dir conditions, at this level.

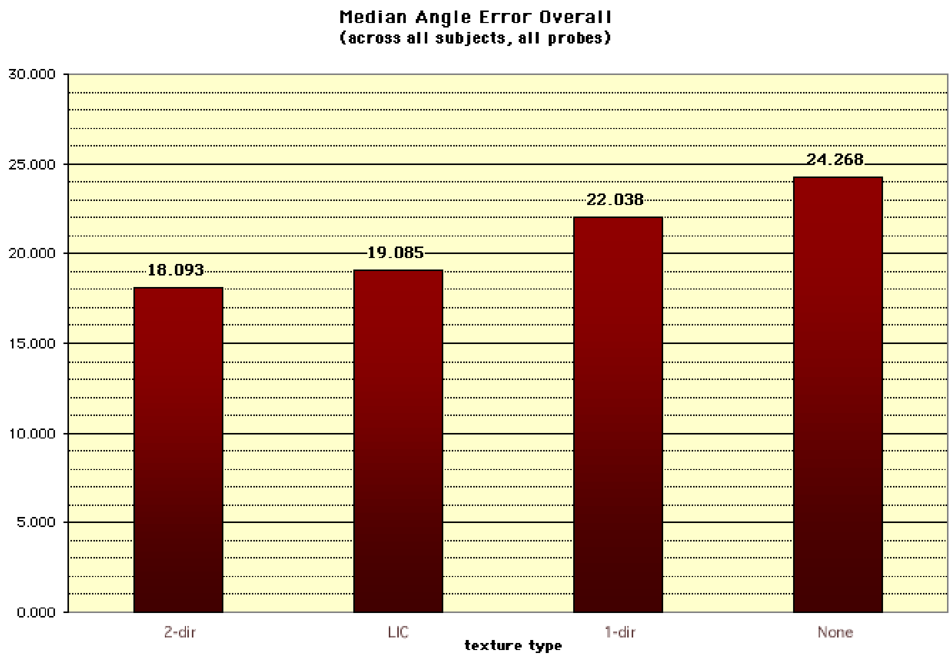

Figure 36. Median angle errors in the different texture conditions, over all subjects and all probe locations.

From the charts in Figure 37 it is possible to gain some deeper insight into the nature of the interaction between probe location and texture type. The first graph shows the median angle errors across all subjects, broken down by probe location. The next 5 graphs show the mean angle errors and standard deviations across the 5 repeated measures for each subject individually, again broken down by probe location.

5.6 DISCUSSION

It appears clear from the results of this experiment that there are small but significant differences in the extent to which various different principal direction-oriented patterns can facilitate accurate shape perception. There are several possible explanations for these results. It could be that performance is better in the case of patterns that carry information about distance along the principal direction than in the case of patterns that only indicate the direction itself. Or it could be that performance is worse in the case of the high contrast, high regularity pattern whose statistics least resemble the statistics of patterns found in nature [23 Párraga]. Despite our concerted efforts to maintain a basic equivalence among the three texture patterns used in this study, it is clear that many differences among the three patterns persist, including differences in spatial frequency and contrast, and differences in higher order statistical characteristics of the patterns. Hence it is possible that the observed performance differences might be alternatively explained by one of these uncontrolled-for differences. For example, there could be an interaction between the spatial frequency of the texture pattern and the resolution accuracy with which surface attitude adjustments can be made, or there could be some interaction between contrast and the availability of shape-from shading information.

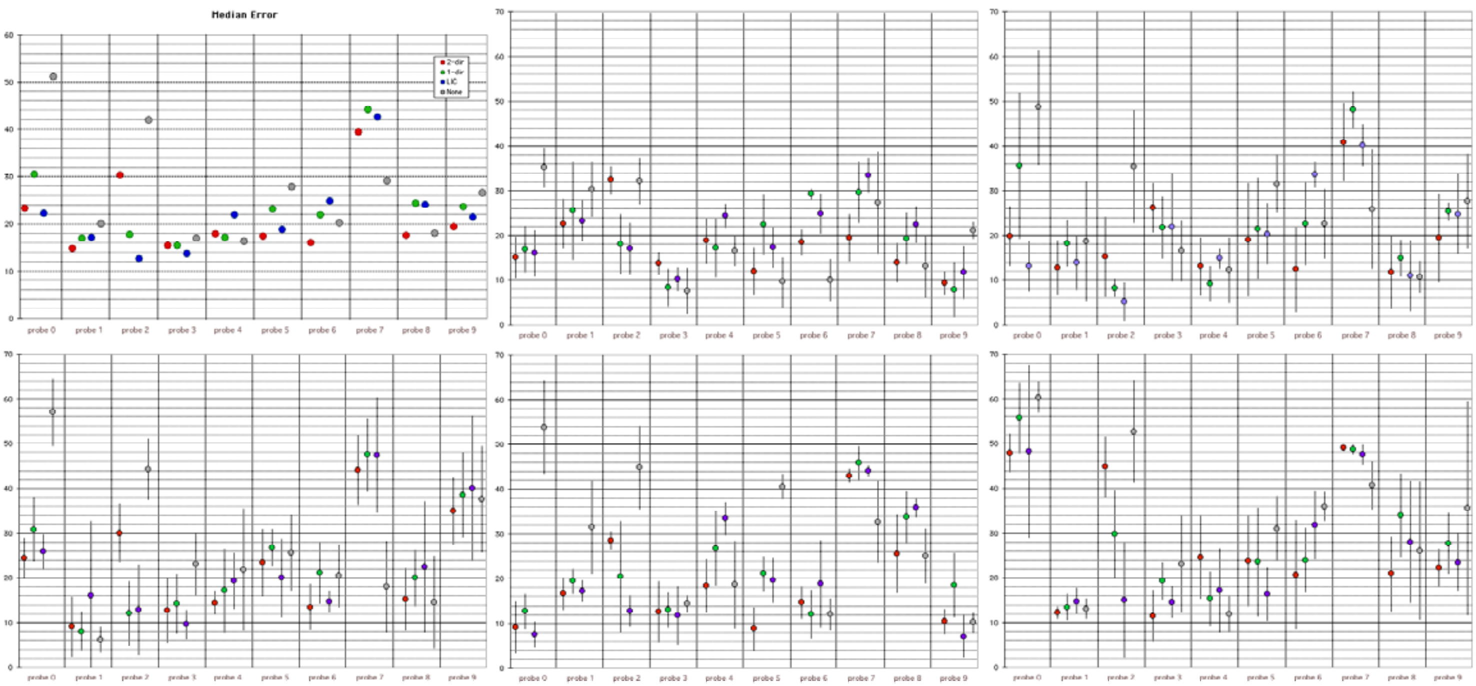

Figure 37. Charts illustrating the details of the experimental results. In the upper left, a graph of the median angle error across all subjects under each texture condition, broken down by probe location. Following that are five graphs showing the mean angle errors and standard deviations for each subject individually, again broken down by probe location. Remarkable consistency can be seen in the pattern of performance by texture type, across subjects, at the same probe locations.

6. SUMMARY

In experiment 1, we examined the effect of the presence and direction of luminance and displacement texture patterns anisotropy on the accuracy of observers' judgments of 3D surface shape. In the monocular viewing condition, performance was significantly better in the cases of the *pdir* and *rdir* patterns than in the cases of the *sdir* and *udir* patterns. In the stereo viewing condition, accuracy increased for all texture types, but was only marginally greater in the cases of the isotropic and principal direction patterns than under the other anisotropic conditions (sinusoidally swirling and uniform direction patterns). These results are consistent with a hypothesis that texture pattern anisotropy can impede surface shape perception when the elongated markings are oriented in a way that is different from the principal direction.

In experiment 2, we attempted to assess the shape information carrying capacity of two different types of directional texture patterns (unidirectional and bi-directional) under three different orientation conditions (following the first principal direction, following a constant uniform direction, or swirling sinusoidally in the surface under two different viewing conditions: frontal and tilted backward). Participants were nearly consistently able to more reliably perceive smaller shape differences when the surfaces were textured with a principal direction texture than when the surfaces were textured with a pattern that either gradually swirled in the surface or followed a constant uniform direction in the tangent plane. There were no apparent significant differences in the pattern of observer responses in the cases of the one-directional vs. two-directional textures (performance was only marginally better overall in the 2-dir case). Also, there was no evidence of an interaction between texture and base surface orientation (tilted vs. front-facing). These findings appeared to support our hypothesis that anisotropic textures aligned with the first principal direction may facilitate shape perception, for a generic view, by making more, reliable information about the extent of the surface curvature explicitly available to the observer than would be available if the texture pattern were oriented in any other way. However they did not yield much insight into the potential effects of principal direction type on shape perception.

In experiment 3, we evaluated the relative effectiveness of three different texture patterns oriented along the principal direction for facilitating accurate shape perception using a surface attitude probe adjustment task. The experiment included a control condition of no texture. Analysis of the results showed that performance was best in the two-directional texture condition, closely followed by the line integral convolution texture that follows first principal direction. Performance was further decreased in the one-directional and no texture conditions (in that order). Tukey's HSD analysis showed that the following differences were statistically significant at the 0.01 level: 2-dir < 1-dir, 2-dir < none, 1-dir < none, and LIC < none. The difference between performance in the 2-dir and LIC conditions was not statistically significant at the 99% confidence level, nor was the performance difference between the LIC and 1-dir conditions.

7. FUTURE WORK

There are exciting opportunities for future work. One of the primary factors motivating this research was the desire to gain insight into how to select or define a texture pattern that could be used to facilitate the accurate and intuitive appreciation of the 3D shape of a rendered surface. Our current work has suggested several possible directions for future investigation:

- *Experiment 1 follow-up.* We tried to replicate the first experiment with both displacement and luminance textures. We hypothesized that: 1) the pattern of errors would be similar to the result of the first experiment; 2) the error would be smaller in displacement texture condition than in luminance texture condition. The results from the replicated experiment confirmed the first hypothesis but we could not get significant results for the second hypothesis. From informal report by observers after they finished the experiment, we suspect that the number of probes observers had to adjust for the experiment was too big and observers may have become very tired towards the end of the experiment decreasing their performance. To limit the complexity of the task in the follow-up, we will randomly choose a small subset of the original probes.

  In the original setup of the experiment, the probes were evenly spaced and this may have led some observers to perceive the probes as being placed on a transparent flat plane in front of the underlying curved surface. Our subjects did not report an inability to see the probes as lying in the surface on an individual basis, but certain of the individual responses appeared to indicate that the probes were not always consistently visualized as a coherent unit across each surface. We plan to address this issue by having the observers view and adjust one probe at a time.

- *Experiment 2 follow-up.* In experiment 2, we investigated the effect of texture orientation over a surface. Ideally, texture orientation should have been the only difference during the trials. However, in our experiment the illumination direction was also different since we wanted to isolate the effect of texture orientation from that of illumination. This may have introduced bias towards "better" lighting direction for same unit differences. While we believe that the results of these experiments are not systematically biased, we would like, for comparison, to obtain new results from a follow-up experiment in which same level differences are textured in all possible texture orientations and the light direction is kept constant. We also plan to add a training phase to the experiment.

- *Experiment 3 follow-up.* In experiment 3, we believe that the amount of global vs. local information necessary for making probe adjustment decisions varied between different texture type conditions and different probe locations. We are in the process of conducting a follow-up experiment in which we explicitly limit the size of the visible area surrounding a probe and compare the results. We hope this will allow us to investigate potential interactions between texture type and task performance as a function of the neighborhood size.

  There seemed to be systematic interaction between probe location and the amount of shape perception error. Some probe locations produced large error for all subjects while in other probe locations the error was smaller for everybody. We would like to investigate this further by analyzing the surface curvature at each probe location.

  Our tool for applying any arbitrary pattern to an arbitrary surface at a high resolution while controlling the pattern orientation at a per-pixel level enables us to pursue investigations of the effects of a wide variety of texture pattern characteristics on shape perception in the more complicated case of doubly curved surfaces. In particular, we have plans to explicitly investigate the impact on shape perception of variations in the contrast and spatial frequency characteristics of a single base pattern (probably LIC, because it is the easiest to control at a fine-tuned level).

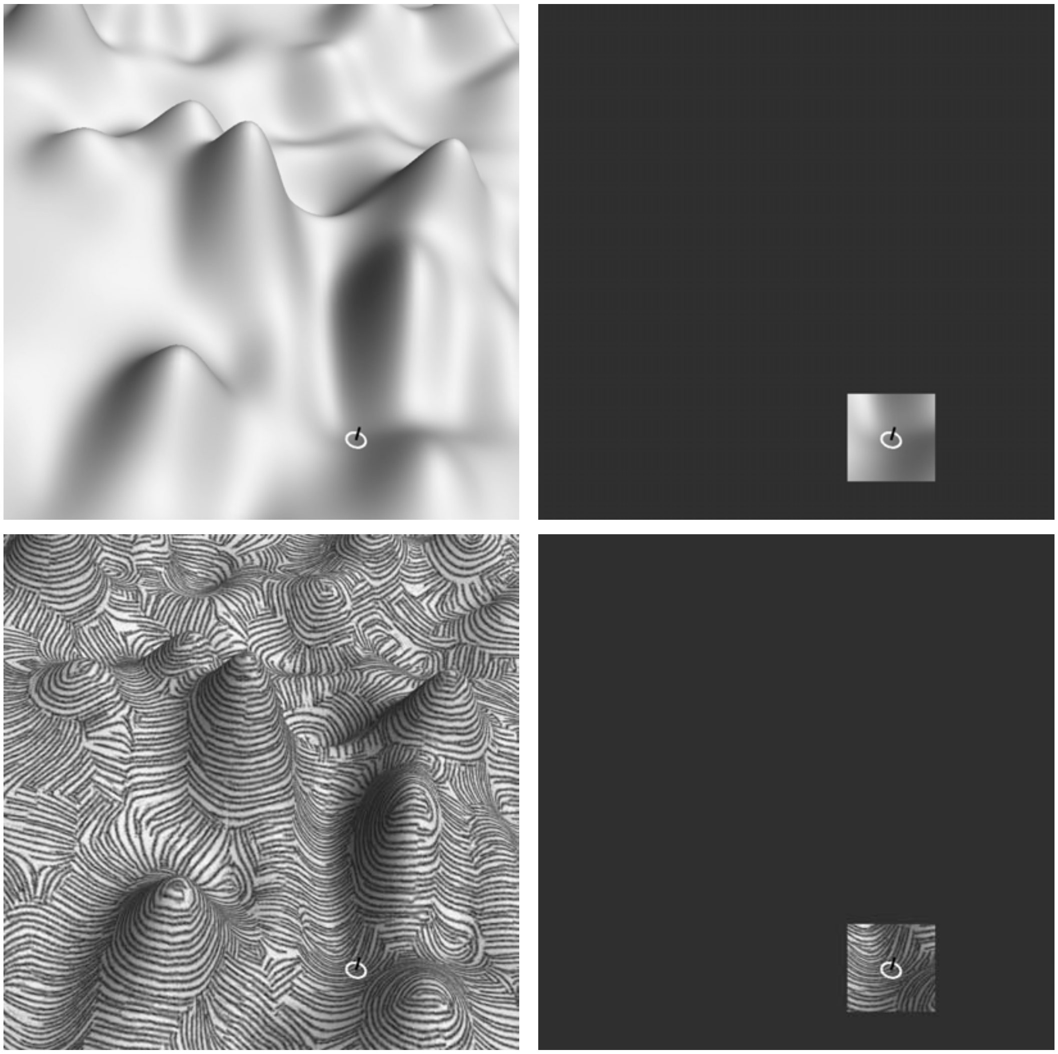

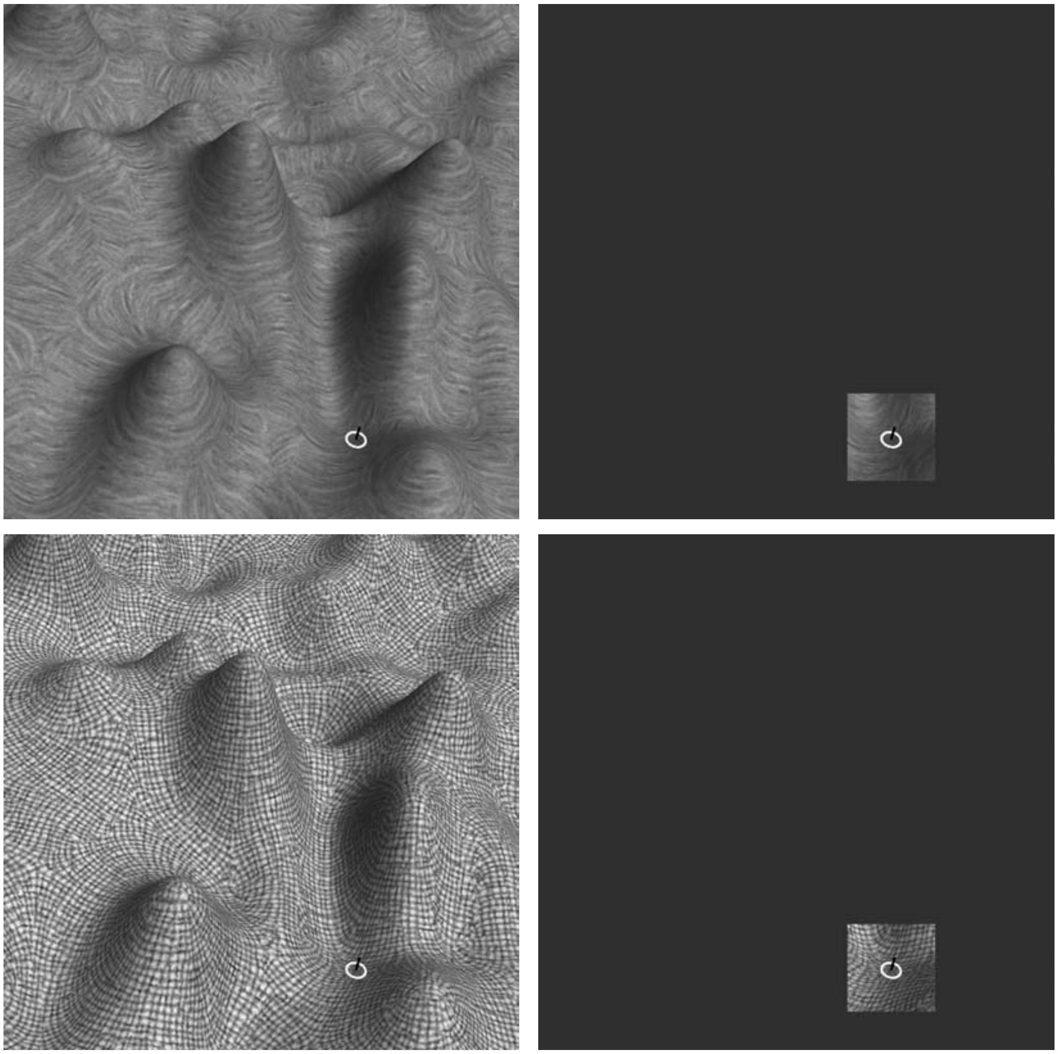

Figure 38. Example stimuli for experiment 3 follow-up: global vs. local information. From top: no texture, 1-dir, LIC, and 2-dir.

- *Experiment 4:* In our studies so far, we displayed our stimuli in perspective projection. In [18, 19], Li and Zaidi argue that: 1) veridical ordinal depth is seen only when the texture contains changes in oriented energy along projected lines of maximum curvature of the surface (first principal direction); 2) texture is a useful cue to shape perception if the image is projected in perspective, since this information is lost in orthographic projection. In related work, Todd and Oomes [28] give examples showing that the perception of surface curvature from texture is only minimally affected by the orientation spectrum of the texture pattern or the amount of perspective in its optical projection. They show that shape-from-texture is possible under both orthographic and perspective projection. We plan to investigate whether perspective projection is necessary for probe-positioning tasks similar to experiment 3. Based on a mock-up experiment we believe that it is not necessary to project the textured surfaces in perspective to perceive local surface curvature. We intend to conduct a formal study with a group of observers and compare our findings with the results from [19, 20, 28].

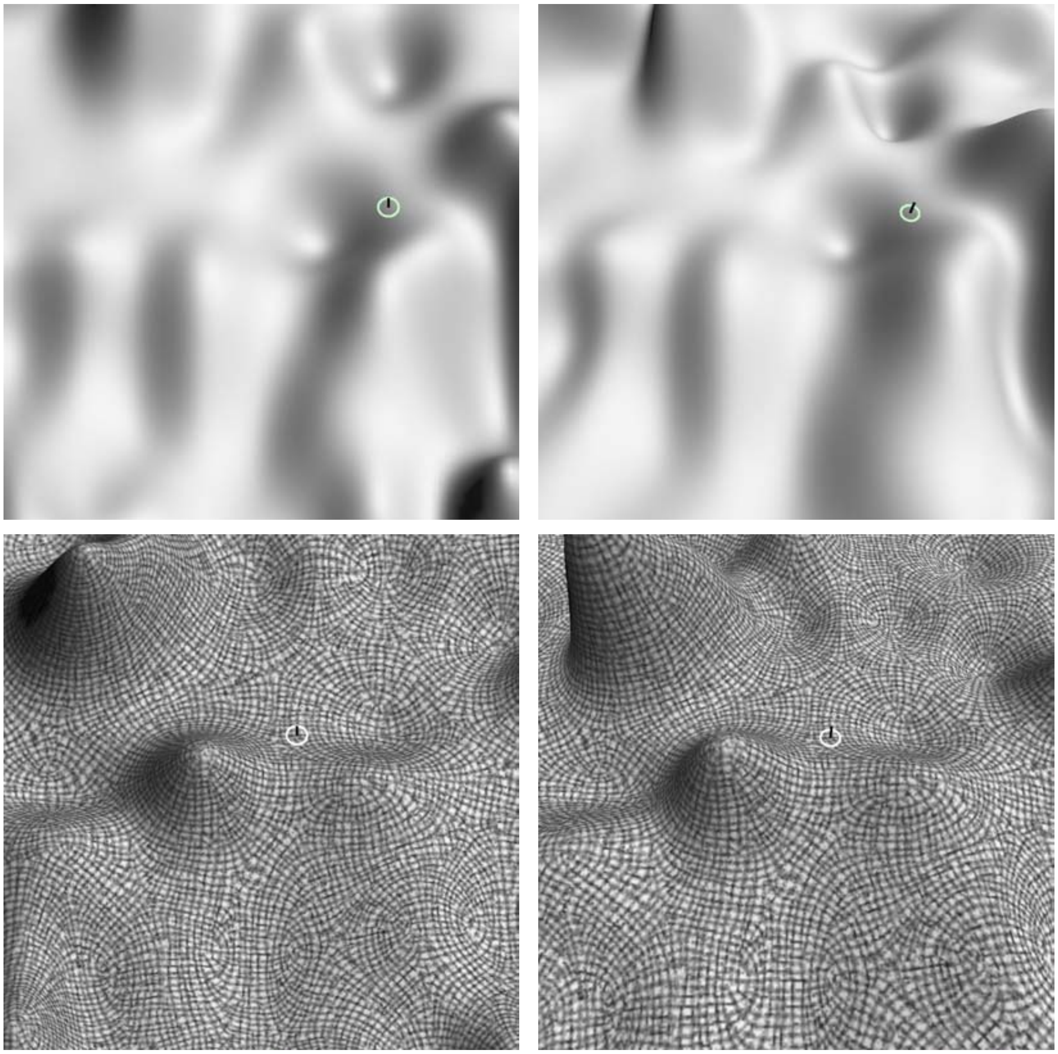

Figure 39. Example stimuli for experiment 4: orthographic vs. perspective projection. Left: surfaces projected in orthographic, Right: surfaces projected in perspective.

- *Experiment 5:* We would like to revisit the question of determining the relative effectiveness of isotropic vs. anisotropic textures for shape representation. In earlier tests we found no significant differences between these two conditions overall, but we suspect that these findings might have been the result of a confluence of several competing factors. In particular, we suspect that certain pattern characteristics, such as the prominent texture flow discontinuities that can arise in 1-directional patterns, are detrimental to shape perception while other aspects, such as the explicit emphasis of the maximal extent of the surface normal curvature in the principal directions, probably facilitate shape perception.

- *Experiment 6:* We believe that some textures show more global shading, which makes direct comparison between texture patterns difficult. For example, LIC texture laid over a surface along the principal direction appears to give both curvature and shading information, while 2-dir texture (Figure 38) gives more of curvature information but may hide shading. We would like to explore this issue by conducting a new experiment in which the amount of global illumination is controlled. We hypothesize that as the amount of global shading is decreased, user performance will deteriorate faster with LIC texture as opposed to 2-dir texture.

Shape representation from line orientation seems to be good in places where one of the two principal curvature values is high, but errors accumulate in the flatter areas where the directional information is less useful and less reliable. Another direction for future work is to develop a more effective texture model that combines the strengths of several different texture definition approaches.

Finally, graphic designers have always been sensitive to the fact that certain patterns are "hard on the eyes" or "annoying to look at", but we are not aware of any formal definition of the characteristics of these patterns, apart from 'extreme regularity'. It is possible that by carefully selecting texture patterns whose statistics match the statistics of natural scenes we can avoid the pitfalls of "annoying" textures and simultaneously improve both the aesthetics and the usefulness of our surface representations and we would like to look into this further.

8. ACKNOWLEDGEMENTS

This work was partially supported by an NSF Presidential Early Career Award for Scientists and Engineers (PECASE). Dr. Julian Rosenman provided the dose data for experiment 1, and the images were rendered using parts of volume rendering software originally written by Marc Levoy. Thanks go to David Banks for suggesting early on that we look into using surface markings that follow both of the principal directions. We are grateful to Paul Schrater and Cindee Madison for many helpful suggestions on the experimental design. We are grateful to Jeremy Leboy for implementing an early version of the surface attitude probe, and to Krista Janssen for developing prototype surface modeling code, with support from the CRA-W Distributed Mentor Project. The texture synthesis and image rendering was done using software originally written by Gabriele Gorla, with support from a University of Minnesota Grant-in-Aid of Research, Scholarship and Artistry, and with co-advice from Guillermo Sapiro. Tim Urness provided us with the LIC texture for experiment 3. Katherine Barnes and Brenda DeBlois provided us with assistance in the experimental analysis through the Statistical Consulting Clinic supported by funding from the Minnesota Agricultural Experiment Station. We are grateful to all our observers MH, TU, GG, MM, and HN in experiment 1, TF and KB in experiment 2, CP, SL, SK, II and TU in experiment 3 for their dedicated and conscientious efforts.

9. REFERENCES

[1] Blake, A. and Marinos, C. "Shape from Texture: Estimation, Isotropy and Moments", *Artificial Intelligence*, **45**, pp. 323-380, 1990.

[2] Brodatz, P. *Textures: A Photographic Album for Artists and Designers*, Dover Publications, 1966.

[3] Biederman, I. "Human Image Understanding: Recent Research and a Theory", in Human and Machine Vision, Azriel Rosenfeld, ed., Academic Press, pp. 13-57, 1985.

[4] Cumming, B. G., Johnston, E. B., and Parker A. J. "Effects of Different Texture Cues on Curved Surfaces Viewed Stereoscopically", *Vision Research,* **33**(5/6), pp. 827-838, 1993.

[5] Ferwerda, J. A., Pattanaik, S. N., Shirley, P., and Greenberg, D. P. "A Model of Visual Masking for Computer Graphics", *ACM SIGGRAPH Proceedings,* pp. 143-152.

[6] Gårding, J. "Shape from Texture and Contour by Weak Isotropy" , *Journal of Artificial Intelligence*, **64**, pp. 243-297, 1993.

[7] Gibson, J. J. "The Perception of Visual Surfaces", *American Journal of Psychology,* **63**, pp. 367-384, 1950.

[8] Gorla, G., Interrante, V., and Sapiro, G. "Texture Synthesis for 3D Shape Representation", *IEEE Transactions on Visualization and Computer Graphics, to appear.*

[9] Interrante, V., Fuchs, H., and Pizer, S. "Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture", *IEEE Computer Graphics and Applications*, **3**(2), pp. 98-117, 1997.

[10] Interrante, V. "Illustrating Smooth Surfaces in Volume Data via Principal Direction-Driven 3D Line Integral Convolution", *ACM SIGGRAPH Proceedings*, pp. 109-116, 1997.

[11] Interrante, V. and Kim, S. “Investigating the Effect of Texture Orientation on Shape Perception”, *Human Vision and Electronic Imaging VI*, SPIE **4299**, pp. 330-339, 2001.

[12] Interrante, V., Kim, S. and Hagh-Shenas, H. “Conveying 3D Shape with Texture: Recent Advances and Experimental Findings”, *Human Vision and Electronic Imaging VII*, SPIE **4662**, pp. 197-206, 2002.

[13] Johnston, A. and Passmore, P. J. “Shape from Shading I: Surface Curvature and Orientation”, *Perception*, **23**, pp. 169-189, 1994.

[14] Knill, D. C. “Contour into Texture: The Information Content of Surface Contours and Texture Flow”, *Journal of the Optical Society of America, A,* **18**(1), pp. 12-35, 2001.

[15] Koenderink, J. J., *Solid Shape*, MIT Press, 1990.

[16] Koenderink, J. J., van Doorn, A. and Kappers A. M. L., “Surface Perception in Pictures”, *Perception*, **52**, pp. 487-496, 1992.

[17] Levoy, M. “Display of Surfaces in Volume Data”, *IEEE Computer Graphics and Applications*, **8**(3), pp. 29-37, 1988.

[18] Li, A. and Zaidi, Q. “Perception of Three-Dimensional Shape from Texture is Based on Patterns of Oriented Energy”, *Vision Research***, 40,** pp. 217-242, 2000.

[19] Li, A. and Zaidi, Q. “Information Limitations in Perception of Shape from Texture”, *Vision Research*, **41**, pp. 1519-1534, 2001.

[20] Lorensen, W. E. and Cline, H. E. “Marching Cubes: A High Resolution 3D Surface Construction Algorithm”, *ACM SIGGRAPH Proceedings*, pp. 163-169, 1988.

[21] Mamassian**,** P. and Landy, M. P. “Observer Biases in the 3D Interpretation of Line Drawings”, *Vision Research*, **38**, pp. 2817-2832, 1998.

[22] Marinos, C. and Blake, A. “Shape from Texture: the Homogeneity Hypothesis”, *Proceedings of ICCV,* pp. 350-353, 1990.

[23] Párraga, C. A, Troscianko, T. and Tolhurst, D. J. “The Human Visual System is Optimized for Processing the Spatial Information in Natural Visual Images”, *Current Biology,* **10**, pp. 35-38, 2000.

[24] Rosenholtz, R. and Malik, J. “Surface Orientation from Texture: Isotropy or Homogeneity (or Both)?”, *Vision Research*, **37**(16), pp. 2283-2293, 1997.

[25] Stalling, D. and Hege, H-C, “Fast and Resolution Independent Line Integral Convolution”, *ACM SIGGRAPH Proceedings,* pp. 249-256, 1995.

[26] Stevens, K. “The Line of Curvature Constraint and the Interpretation of 3-D Shape from Parallel Surface Contours”, *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 1057-1061, 1983.

[27] Todd, J. T. and Reichel, F. D. “Visual Perception of Smoothly Curved Surfaces from Double-Projected Contour Patterns”, *Journal of Experimental Psychology: Human Perception and Performance*, **16**(3), pp. 665-674, 1990.

[28] Todd, J. T. and Oomes A. H. J. “Generic and non-generic conditions for the perception of the surface shape from texture”, *Vision Research,* **42**, pp. 837-850, 2002.

[29] Witkin, A. P. “Recovering Surface Shape and Orientation from Texture”, *Journal of Artificial Intelligence*, **17**, pp. 17-45, 1981.

[30] Zaidi, Q. and Li, A. “Limitations on shape information provided by texture cues”, *Vision Research,* **42**, pp. 815-835, 2002.