Investigating the Effect of Texture Orientation on the Perception of 3D Shape

Victoria Interrante and Sunghee Kim

Department of Computer Science and Engineering
University of Minnesota

ABSTRACT

Perception of the 3D shape of a smoothly curving surface can be facilitated or impeded by the use of different surface texture patterns. In this paper we report the results of a series of experiments intended to provide insight into how to select or design an appropriate texture for shape representation in computer graphics. In these experiments, we examine the effect of the presence and direction of luminance texture pattern anisotropy on the accuracy of observers' judgments of 3D surface shape. Our stimuli consist of complicated, smoothly curving level surfaces from a typical volumetric dataset, across which we have generated four different texture patterns via 3D line integral convolution: one isotropic and three anisotropic, with the anisotropic patterns oriented across the surface either in a single uniform direction, in a coherently varying direction, or in the first principal direction at every surface point. Observers indicated shape judgements via manipulating an array of local probes so that their circular bases appeared to lie in the tangent plane to the surface at the probe's center, and the perpendicular extensions appeared to point in the direction of the local surface normal. Stimuli were displayed as binocularly viewed flat images in the first trials, and in stereo during the second trials. 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 than in the cases of the other two anisotropic pattems. Under stereo viewing, accuracy increased for all texture types, but was still greater for the isotropic and principal direction patterns than for the other two. Our results are consistent with a hypothesis that texture pattern anisotropy impedes surface shape perception in the case that the direction of the anisotropy does not locally follow the direction of greatest normal curvature.

Keywords: texture, shape representation, principal directions, shape perception.

1 INTRODUCTION

A key objective in the field of visualization is to design and implement algorithms for effectively communicating information through images. Given a set of data, we must design a visual representation for that data which facilitates its understanding. The investigations reported in this paper were motivated by applications in which one needs to be able to accurately and intuitively convey the three-dimensional shape of large, smooth, arbitrarily curving surfaces. Previous studies [cf. Interrante *et al.* 97] have indicated that shape perception can facilitated by the addition of surface texture markings, but the question of how to characterize the kind of surface texture that will show shape best remains open. Texture can also be used to *mask* surface shape features, as was shown by Ferwerda *et al.* [1997]. It has been suggested [Cumming *et al.* 1993] that the perception of shape from texture can be impeded when the texture pattern is highly anisotropic, consisting of elements that are systematically elongated in a specific direction. However a wide variety of textures consisting of line-like elements have been shown to indicate surface curvature [Todd and Reichel 1990]. Knill [1999] shows that across developable surfaces[1] any homogeneous texture pattern will appear to flow along parallel geodesics[2], and suggests that our visual system uses shape-from-contour[3] to infer shape from the systematic projective distortion or flow of the pattern. Inspired by the considerable amount of research [Stevens 1983, Mamassian and Landy 1998, Li and Zaidi 2000] that seems to imply that surface shape may be perceived most accurately from line-like markings when they follow the lines of curvature., we sought in the series of experiments described in this paper to further experimentally investigate the effect of the direction of surface texture pattern anisotropy on the accuracy of observers' shape judgments. Specifically we wanted to know: 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? 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?

Correspondence: Interrante; Other author info: VI; email: interran@cs.umn.edu; web: www.cs.umn.edu/~interran; SK; email: skim@cs.umn.edu; web: www.cs.umn.edu/~skim.

[1] surfaces that can be unrolled to lay out flat in the plane;

[2] curves that do not turn in the surface

[3] the set of all points at which the surface normal is orthogonal to the line of sight

2 METHODS

We conducted a series of two experiments intended to investigate the effect of the presence and direction of texture pattern anisotropy on the ability of observers to accurately perceive the 3D shape of a smoothly curving surface. The goal of these experiments was to gain insight that might facilitate our efforts to use texture most effectively to facilitate the accurate perception of surface shape in renderings of scientific data. In the following sections we provide the details of the experimental set up and design.

2.1 Stimuli

The stimuli that we used in our experiments were cropped images of the front-facing portions of textured level surfaces rendered in perspective projection using a hybrid renderer [Interrante *et al.* 97] that uses raycasting [Levoy 88] together with a Marching Cubes algorithm [Lorensen and Cline 87] 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 testbed, rather than a more restricted type of analytically-defined surface, because this data is typical of the kind of data whose shape features we seek to be able to more effectively portray through the use of surface texture.

The first step in image generation was to define the solid texture patterns that would appear on the level surfaces. We used a high-quality three-dimensional line integral convolution algorithm [Stalling and Hege 95] to synthesis the 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 [Interrante 97], but is briefly restated here for completeness. We begin by computing an orthogonal frame at each sample point in the 433x357x325 voxel 3D volumetric dataset. We define the third frame vector to be in the direction of the grey-level gradient, which is the normal to the level surface that passes through the sample. We compute the gradient using Gaussian-weighted central differences in the axial directions over the 3x3x3 area surrounding the sample point. We next choose an arbitrary point in the tangent plane to define the direction of the first frame vector and then take the cross product of these two vectors to obtain the remaining orthogonal direction. Finally, we estimate the 2nd Fundamental Form [Koenderink 90] 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 (eigenvectors) 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 obtained by simpler means. First, we obtain the vector field of uniform directions 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: $udir_x = -n_y$, $udir_y = n_x$, $udir_z = 0$, where (n_x, n_y, n_z) is the surface normal or gradient. Then, we obtain the vector field of random directions 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, $-\pi/2 \leq \theta_1 \leq \pi/2$. Finally, we obtain the vector field of coherently varying directions that is 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.

Figure 1 illustrates the process of texture synthesis, showing a single slice (z=263) from the 3D input texture volume and from each of the different 3D output texture volumes. Not all of the values are filled in, because we have elected to initiate the streamlines that are used to compute the output texture values only at the voxels that are in the vicinity of the level surface that is being used to create the test image.

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, and Phong shading is then 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 2 shows three of these images, all computed for the same viewing position. In order to avoid the potentially confounding influence of shape-from-contour information, as a last step we cropped each image to a 400x400 pixel region that did not contain any points on the silhouette edges of the object.

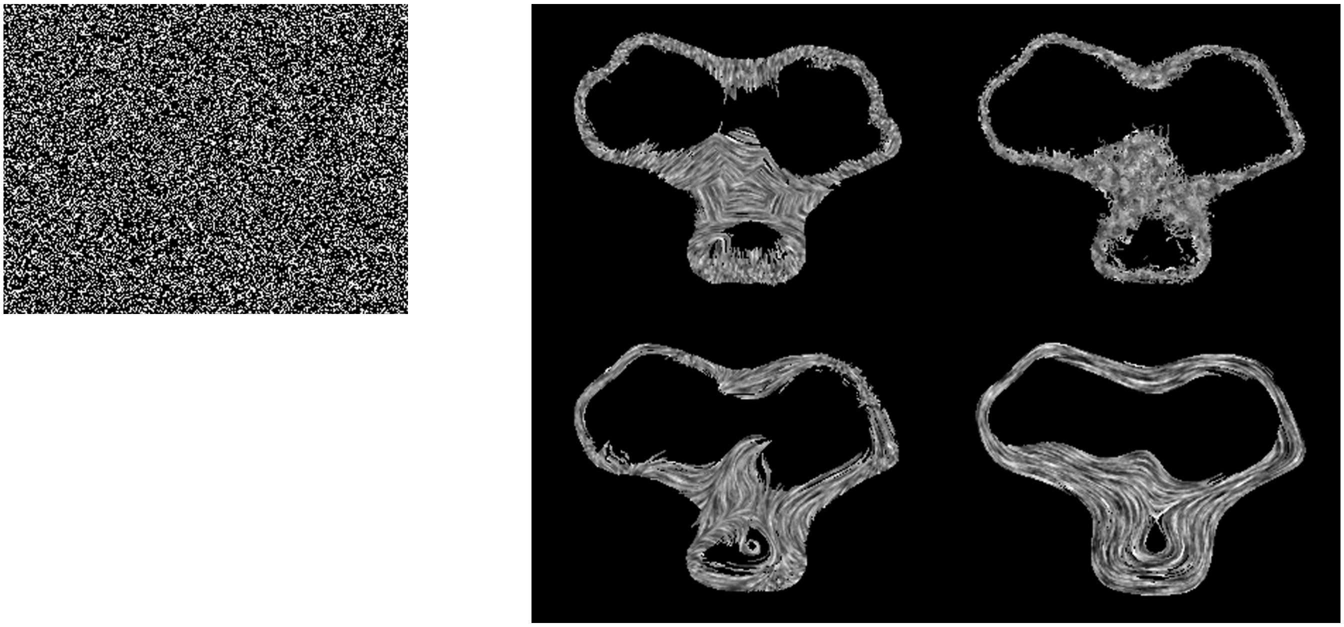

Figure 1: Slices from the 3D solid textures. Left: The slice z=263 before line integral convolution; Right: The same slice after line integral convolution along (in clockwise order) first principal directions (pdir), random directions (rdir), uniform directions (udir) and coherently varying or swirling directions (sdir) computed at each sample point.

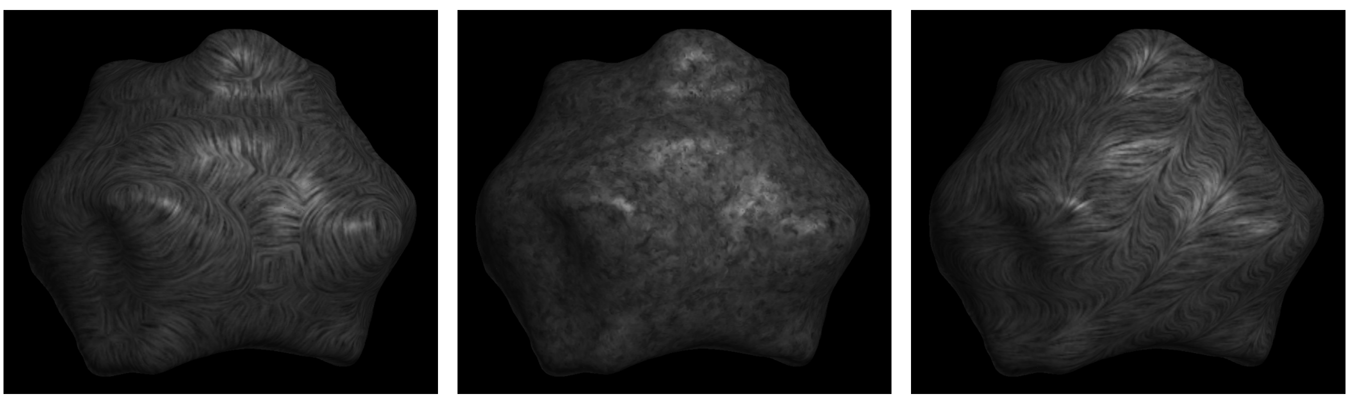

Figure 2: Examples of the 3D textured surfaces. From left to right: pdir, rdir and sdir. Note that informal assessment of the potential impact of texture type on shape judgments is complicated in these images by the prominence of shape-from-contour cues, which tend to dominate when other information about shape is less readily accessible. Because we are most interested in studying how the presence of texture might facilitate shape judgments across non-trivially structured interior regions where shape-from-contour information is not available, we cropped all of the images to eliminate the edge cues before testing.

2.2 Task

In originally planning these investigations, we had hoped to be able 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). However we had great difficulty coming up with a means to evaluate observers' immediate global impressions of surface shape in a way that avoided confounding influences such as isolated 2D feature recognition or partial picture matching. Hence we decided to proceed with estimates of surface shape perception accumulated from individual judgments of the orientation of the surface at local points. Because 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, but rather obtains shape understanding from the comparative relationships between nearby points, we decided to present an array of probes [Koenderink *et al.* 1992] that completely covered the central area of the presented surface and to ask observers not only to adjust each probe by pulling on its handle until the circular base appeared to lie in the tangent plane to the surface at its central point and the perpendicular extension appeared to point in the surface normal direction, but also before proceeding to the next trial to verify that the shape of the surface they had implicitly indicated through the collective orientations of all of the probes appeared to faithfully match the shape of the underlying textured surface at all points.

Unfortunately, we neglected to recognize, before beginning the experiments, that our decision to place the probes at exactly evenly spaced intervals over a rectangular grid would interfere with observers' ability 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 arrayed 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, as will be discussed later, certain of the individual responses appeared to indicate that the probes were not always consistently visualized as a coherent unit across each image. Figure 3 shows the user interface at the beginning of the 5th trial.

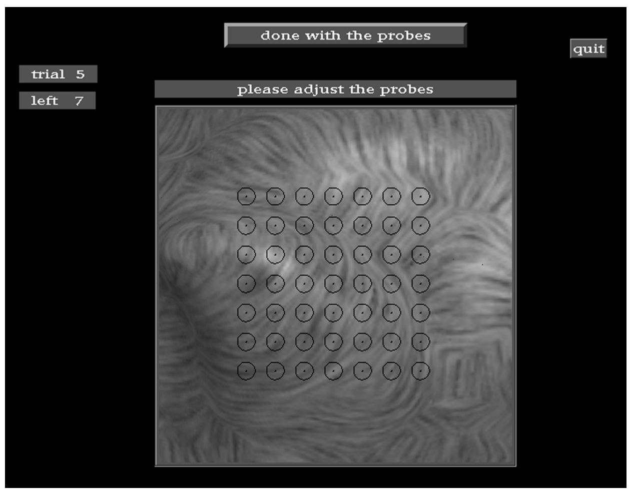

Figure 3: The graphical user interface with all probes displayed in their starting positions.

In designing the experiment we were particularly concerned about avoiding a situation in which differences dues to texture type might be confounded with differences due to other unanticipated or uncontrolled factors such as individual differences, or particular surface shape configurations. Ideally, we would have liked to present identical views of each surface under all four texture conditions, and to have all subjects make judgements on all of the images. However we were also concerned about the possibility of subjects' current shape judgments being biased by information they obtained from previous trials in which the same surface had been shown under a different texture condition. With only 24 binocular images (6 views x 4 texture types), and an small anticipated subject pool size, we had to make some difficult tradeoffs. What we did was to divide the subjects and the stimuli into two different groups, so that each subject made shape judgments at the 49 probe locations on only half of the data (12 images). Each set contained each view and each texture type, in equal proportions, but did not contain all of the possible combinations. Within each set, the stimuli were further grouped into two lots, in which each lot had no surface repeated. Figures 4 and 5 show the complete set of stimuli presented to each group of observers.

The six images within each lot were presented in random order, and subjects were required to take a 10 minute rest break after finishing the 6th trial, thereby avoiding the possibility that any two differently textured but identically shaped surfaces might be presented immediately in sequence, and minimizing the likelihood of surface recognition and any consequent possible learning effects. After adjusting the probes on the 12 images in the flat viewing condition, subjects repeated the entire process under conditions of stereo viewing. To facilitate estimation of the effects of viewing condition, subjects were presented with the same stimuli, differently ordered, in the two viewing conditions. The entire process took about two hours for most of the subjects.

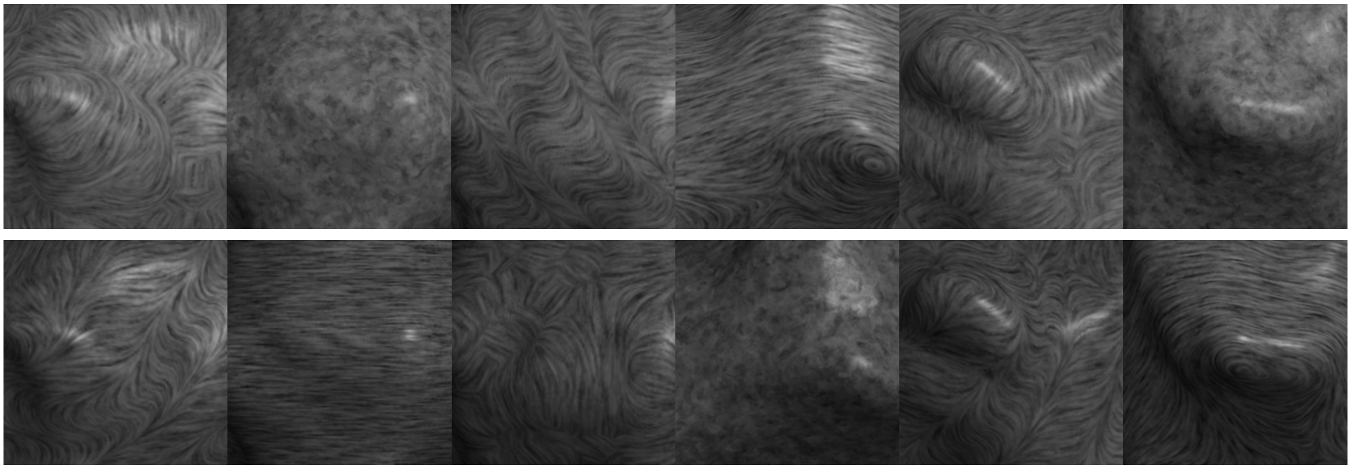

Figure 4: The set of stimuli seen by group A. First row: lot 1; second row: lot 2. The presentation order was randomly determined and was different for each subject.

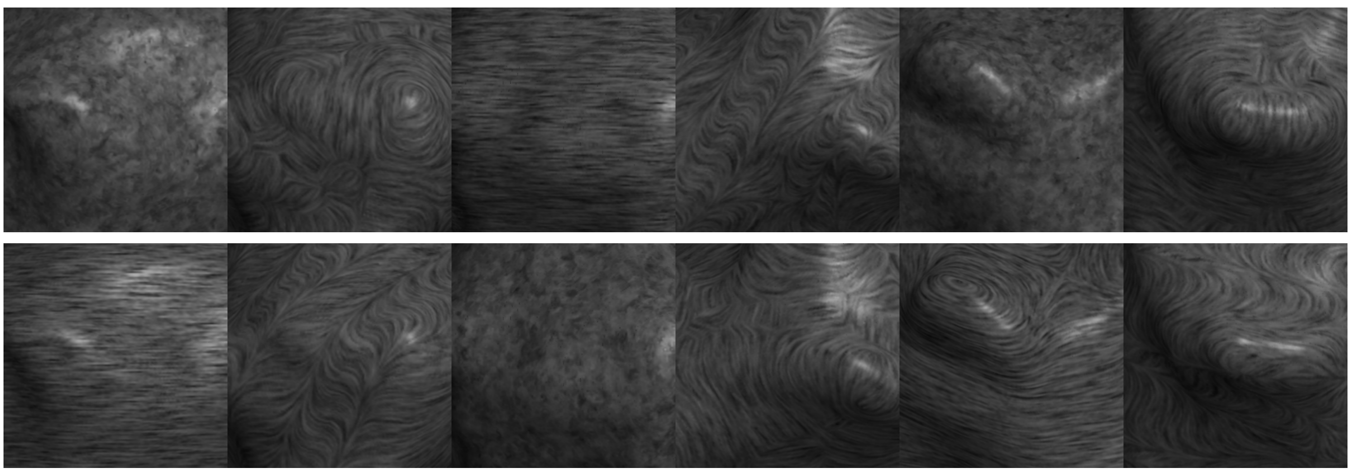

Figure 5: The set of stimuli seen by group B. First row: lot 1; second row: lot 2. The presentation order was randomly determined and was different for each subject.

2.3 Observers

We had five subjects participate in the experiments. All of the subjects were male EE and CS graduate students from the University of Minnesota, who agreed to participate as a favor to the second author and for compensation in the form of gift certificates to local coffee shops and/or eateries. All subjects were kept fairly naïve to the purposes of this experiment, though some of the subjects were certainly aware of the authors' previous work with principal direction texture. We informed the subjects that we were conducting experiments to evaluate peoples' ability to accurately perceive 3D shape in images but we specifically did not mention anything about texture. Our goal in doing this was to keep the subjects as free as possible of any potential biases and to avoid leading them into certain behaviors (such as lining up the direction of probe base elongation with the direction of the texture pattern) that they might not otherwise have considered. Before beginning the experiment, the subjects were asked to read a set of written instructions which described the probe positioning task. We used written rather than verbal instructions in an effort to maintain consistency. Subjects were also shown a single "training" image (figure 6) that portrayed ground truth answers in the form of correctly positioned probes for a seventh surface not included in the test data and rendered without texture. Note that several of these probes appear to point straight out of the screen. We showed them this image in order to give them an idea of what a set of exactly correctly positioned probes might look like. We were fairly selective in attempting to obtain participants that we hoped would be diligent in their efforts, and in the written instructions we stressed the importance of trying hard to do a consistently good job on all of the images, even if the shape was difficult to perceive. As an extra incentive, however, we told the subjects that after all results had been tallied, we would give a $20 bonus certificate to the student who gave the most accurate answers, overall.

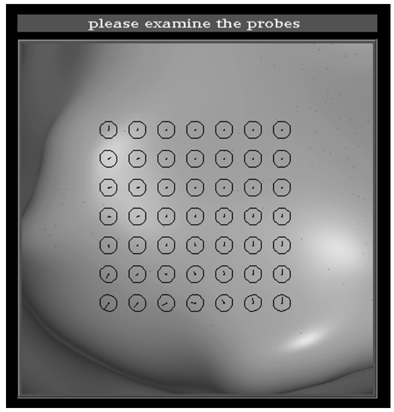

Figure 6: Training image, showing ground truth answers (correct probe orientations) at points across an untextured surface.

3 RESULTS

Having observed that other investigators studying shape perception using local probes analyze the perceived surface orientation in terms of slant and tilt, where slant is the angle of rotation out of the fronto-parallel plane, and the tilt is the angle of rotation about the viewing direction, we had initially hoped to be able to do the same in our studies – measuring the accuracy of observers' estimates of local heading in terms of the deviation in slant and tilt from the ground truth answers. While fairly satisfied with the indication of error provided by deviations in slant, we had several serious problems interpreting the magnitude of the errors due to incorrect estimates of tilt. The root of our difficulties was that too many of the points on our surfaces were too near to being parallel with the image plane. In numerous incidences the angular deviation in tilt was degenerate, because the estimated normal projected to a single point, and it was not clear how to appropriately handle these cases. 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, where the cues to shape were inadequate and subjects reverted to the default assumption that the surface lay in the plane of the image, or subjects simply gave up in frustration and left the probes untouched at their default original positions. Furthermore, even in the cases where the tilt angle was not degenerate, the lengths of the projected normal vectors could be exceedingly small, on the order of one or two pixels, and it was therefore possible to register huge estimated errors in the tilt component in places where the observer had merely misplaced the endpoint of the vector by two or three pixels (less than 1mm on the screen) in a particularly unfortunate direction. We therefore reluctantly decided to break with tradition and simply use 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 7 shows the mean angular error and standard deviations computed over the 49 probe positions at which estimates were made by each subject for each image, with each texture type, under conditions of binocular flat viewing. The results are grouped into different images by texture type, and then grouped within each image by test subject. Figure 8 shows the results under conditions of stereo viewing.

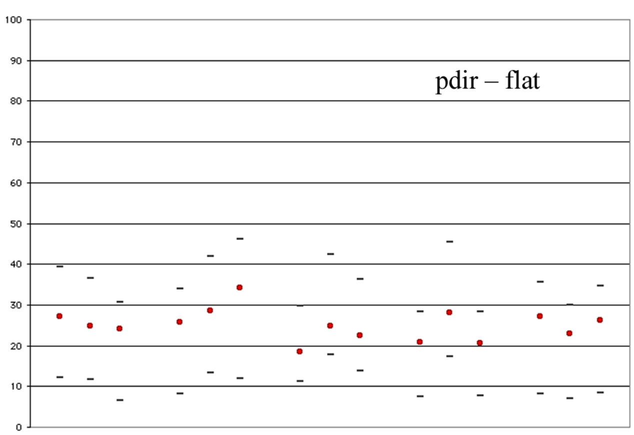

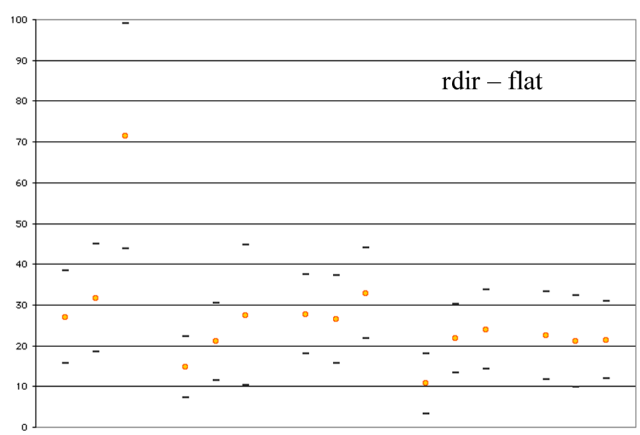

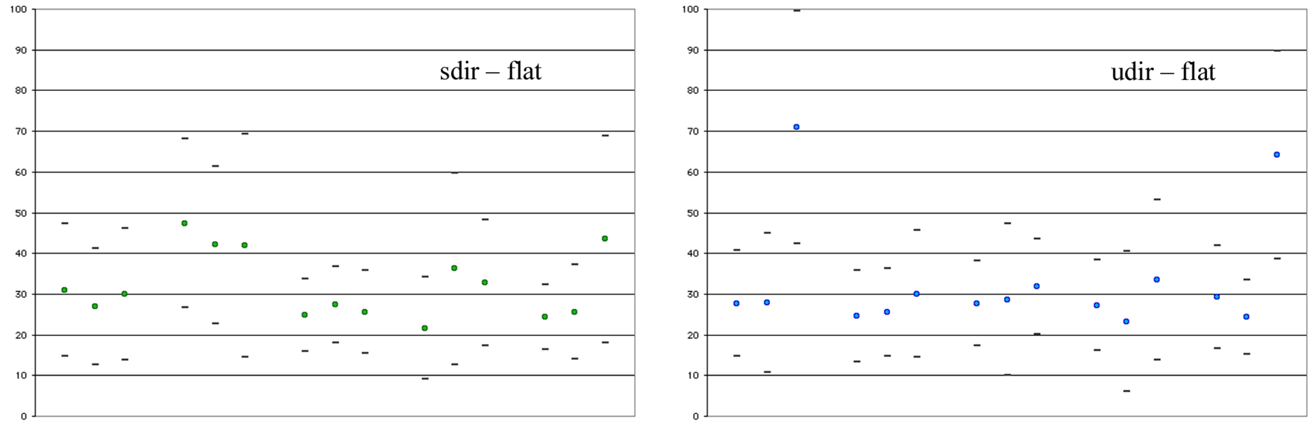

Figure 7: Individual results for the flat viewing condition. The height of each point represents mean angular error over the 49 probe locations per image. Subject number is the unspecified independent variable along the horizontal axis. Judgements from a single subject for different surfaces rendered 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 swirling (sdir).

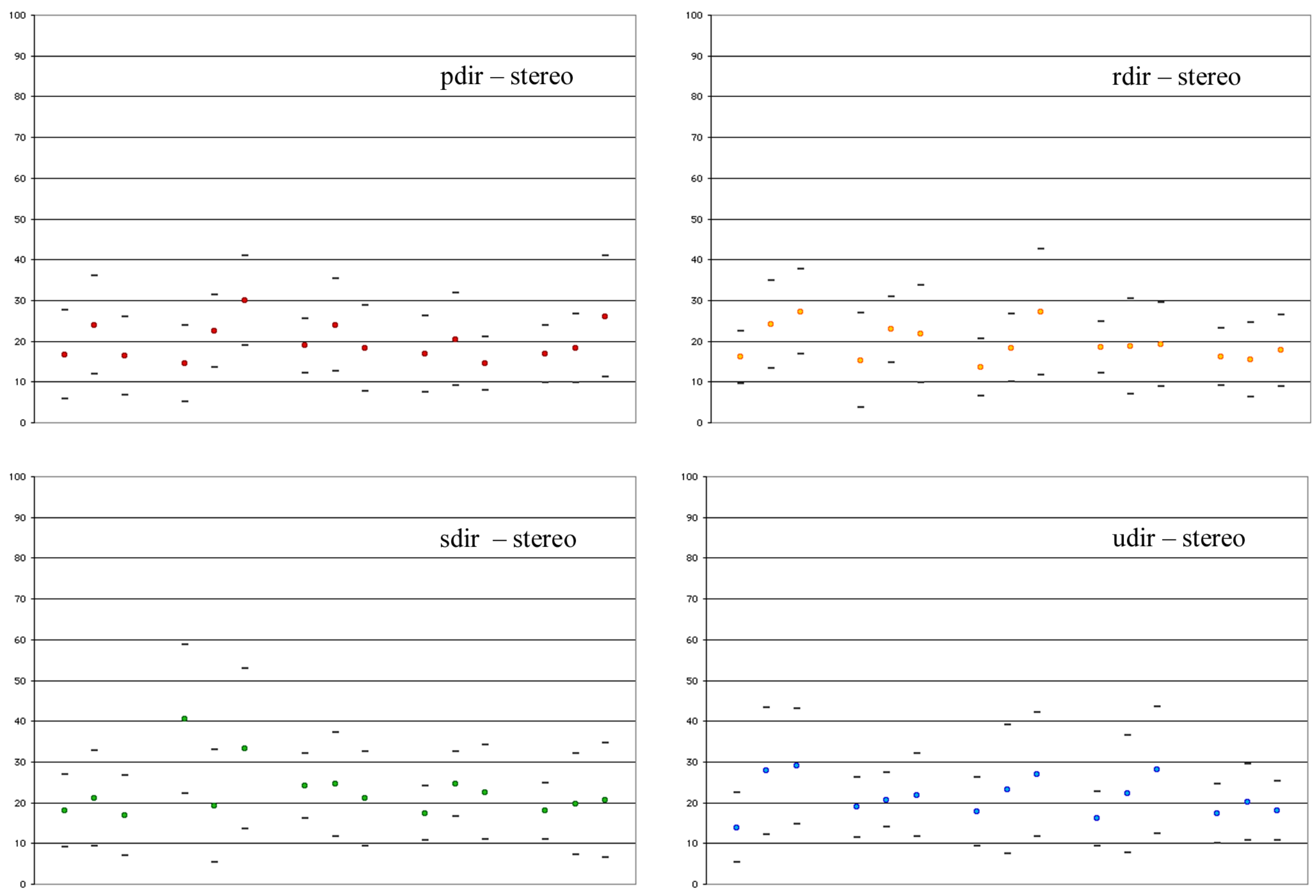

Figure 8: Individual results for the stereo viewing condition. Each point represents mean angular error over the 49 probe locations per image. Clockwise from top left: principal direction (pdir), isotropic (rdir), uniform (udir), swirling (sdir).

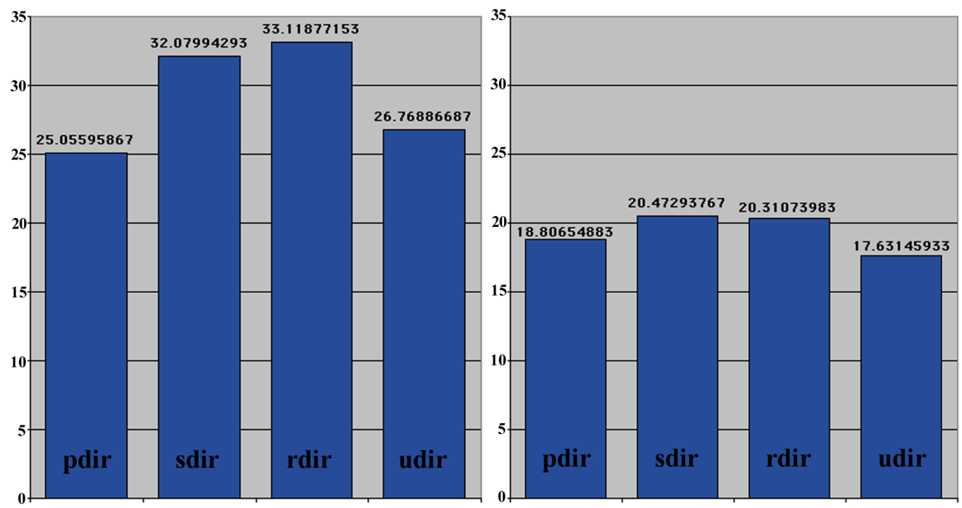

Figure 9: Pooled results (mean angle error) for all subjects, all surfaces, by texture type. Left: flat presentation; Right stereo.

4 DISCUSSION

A definitive statement of the results is hampered by the fact that we have not yet succeeded in doing a thorough statistical analysis of the data and hence cannot make any claims about the statistical significance of the differences in the mean angular errors observed under the different texture conditions. Overall, subjects seem to do somewhat better in the principal direction oriented and isotropic texture conditions than in either of the other two. It appears, from inspection of the individual results, that subjects may be less prone to making catastrophic errors when stimuli are viewed as flat images if the surfaces are rendered with the pdir texture. However, closer inspection of the pattern of errors is needed. A preliminary inspection suggests the presence of two different types of errors: coherent errors due to perceived depth inversion, and incoherent errors, as shown in figure 10. 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. Most subjects appeared to perform equally well or better with the purely isotropic pattern. However some subjects were clearly misled in some places by the anisotropic patterns that followed directions different from the principal direction, suggesting that if one must use an anisotropic pattern, one must be careful about how it is applied over the object.

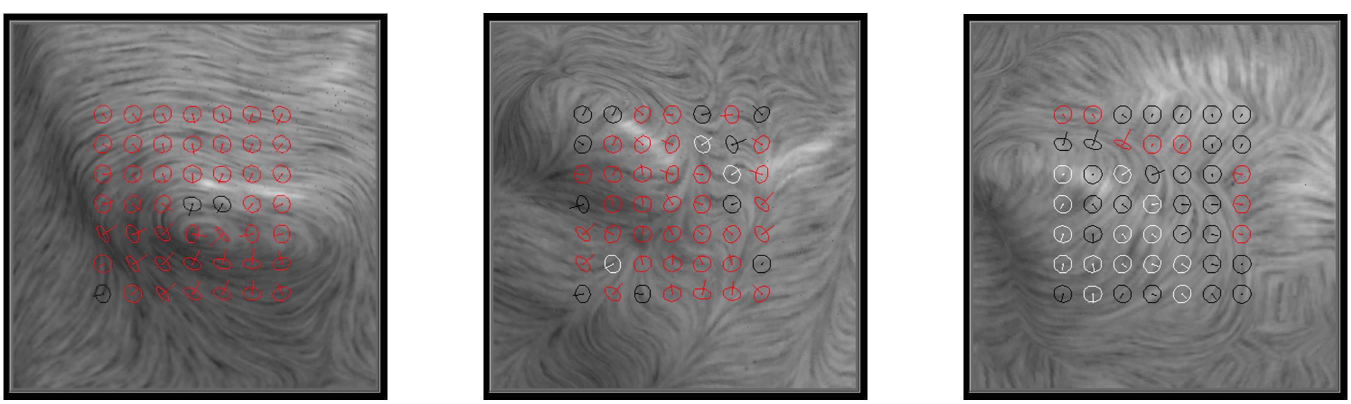

Figure 10: Some detailed individual results: Left: coherent errors due to depth inversion; Middle: incoherent errors apparently due to shape misperception; Right: errors tend to pile up at texture flow discontinuities, where the first and second principal directions switch places.

5 FUTURE WORK

There is considerable room 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 3D shape of a rendered surface. It appears clear that the principal direction textures defined above leave something to be desired in this respect. 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. One direction for future work is to develop a more effective texture model that combines the strengths of several different texture definition approaches. A perhaps more immediate direction for future work is the investigation of the effect of texture orientation on surface shape judgments when the texture pattern is defined by surface relief rather than surface luminance. Does texture orientation affect shape perception in the same way in the two cases? Examples of some preliminary stimuli for subsequent experiments on this subject are shown in figure 11.

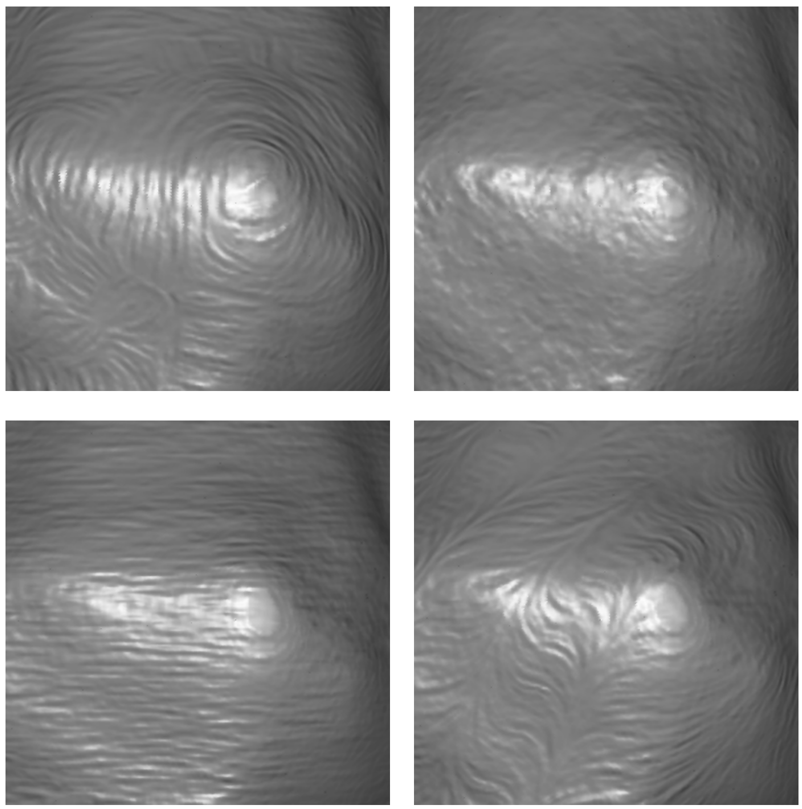

Figure 11: The same stimuli with the same textures, this time rendered as shaded relief rather than as luminance patterns.

6 ACKNOWLEDGMENTS

This work was partially supported by an NSF Presidential Early Career Award for Scientists and Engineers (PECASE). The images were rendered using parts of volume rendering software originally written by Marc Levoy. The dose data was provided by Dr. Julian Rosenman. Jeremy Leboy assisted with the implementation of the probe. We are grateful to our subjects TU, GG, MH, MM and HN for their dedicated efforts.

7 REFERENCES

Cumming, Bruce G., Elizabeth B. Johnston and Andrew J. Parker. (1993) Effects of Different Texture Cues on Curved Surfaces Viewed Stereoscopically, *Vision Research*, **33**(5/6): 827–838.

Ferwerda, James A., Sumanta N. Pattanaik, Peter Shirley and Donald P. Greenberg. (1997) A Model of Visual Masking for Computer Graphics, *ACM SIGGRAPH Proceedings*, pp.143–152.

Interrante, Victoria, Henry Fuchs and Stephen Pizer. (1997) Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture, *IEEE Computer Graphics and Applications*, **3**(2): 98–117.

Interrante, Victoria. (1997) Illustrating Smooth Surfaces in Volume Data via Principal Direction-Driven 3D Line Integral Convolution, *ACM SIGGRAPH Proceedings*, pp. 109–116.

Knill, David C. (1999) Contour into Texture: The Information Content of Surface Contours and Texture Flow, *www.cvs.rochester.edu/~knill/papers/postscript/texture_flow.pdf.*

Koenderink, Jan J. (1990) Solid Shape, MIT Press.

Koenderink, Jan J., Andrea van Doorn and Astrid M. L. Kappers. (1992) Surface Perception in Pictures, *Perception*, **52**, pp. 487–496.

Levoy, Marc. (1988) Display of Surfaces in Volume Data. *IEEE Computer Graphics and Applications*, **8**(3): 29–37.

Li, Andrea and Qasim Zaidi. (2000) Perception of Three-Dimensional Shape from Texture is Based on Patterns of Oriented Energy, *Vision Research*, **40**, pp. 217–242.

Lorensen, William E. and Harvey E. Cline (1987) Marching Cubes: A High Resolution 3D Surface Construction Algorithm, *ACM SIGGRAPH Proceedings*, pp.163–169.

Mamassian, Pascal and Michael P. Landy. (1998) Observer Biases in the 3D Interpretation of Line Drawings, *Vision Research*, **38**, pp. 2817–2832.

Stalling, Detlev and Hans-Christian Hege. (1995) Fast and Resolution Independent Line Integral Convolution. *ACM SIGGRAPH Proceedings*, pp. 249–256.

Stevens, Kent (1983) 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.

Todd, James T. and Francene D. Reichel. (1990) The Visual Perception of Smoothly Curved Surfaces from Double Projected Contour Patterns, *Journal of Experimental Psychology: Human Perception and Performance*, **16**(3): 665-674.