Showing shape with texture – two directions seem better than one

Sunghee Kim[1], Haleh Hagh-Shenas[1], Victoria Interrante[1]
Dept. of Computer Science and Engineering, University of Minnesota, Minneapolis, MN

ABSTRACT

Studies have shown that observers' judgments of surface orientation and curvature are affected by the presence of surface texture pattern. However, the question of designing a texture pattern that does not hide the surface information nor conveys a misrepresentation of the surface remains unsolved. The answer to this question has important potential impact across a wide range of visualization applications. Molecular modeling and radiation therapy are among the many fields that are in need of accurately visualizing their data that could benefit from such methods.

Over the past several years we have carried out a series of experiments to investigate the impact of various texture pattern characteristics on shape perception. In this paper we report the results of our most recent study. The task in this study was adjusting surface attitude probes under three different texture conditions and under a control condition in which no texture was present. The three texture conditions were: a doubly oriented texture in which approximately evenly spaced lines followed both of the principal directions, a singly oriented texture in which lines followed only the first principal direction, and a singly oriented line integral convolution like texture. In a series of 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 surface attitude probe. Probes were randomly positioned on one of five different arbitrary curved surfaces and the observer's task was to adjust the orientation of the probe so that its base appeared to lie in the surface and its perpendicular extension appeared to be oriented in the direction of the surface normal. An analysis of the results showed that the performance was best in the two directional grid texture and like-like texture conditions. Performances were significantly decreased in one directional line texture and no texture conditions (in that order). The paper is organized as follows. In Section 1 we briefly describe the motivation for our work. In Section 2 we describe our experimental methods, including a brief summary of the process of the stimuli preparation and a detailed presentation of the statistical analysis of our experimental results. In Section 3 we discuss the implications of our findings and in the Section 4 we talk about our future plans.

Keywords: Shape representation, shape perception, texture synthesis, texture mapping.

1. INTRODUCTION

As visualization designers, our goal is to determine the most effective way of portraying a set of data so that its essential features can be understood easily and accurately. When we choose to render a surface or an object, we have tremendous latitude in choosing how we want to model its material properties. The most common practice is to use a simple Phong shading model without any surface texture. This model is commonly used because it is easy to implement and is the default model on most systems. However it is becoming increasingly clear that this model is not 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 tell us how to specify the surface material properties of an object in order to best facilitate the accurate understanding of its shape. In broad terms, the fundamental objective of the study reported in this paper is to seek further quantitative insights the effects of texture pattern characteristics on surface shape perception, better enabling informed visualization design decisions by helping to build up an solid theoretical foundation for knowing what works, what doesn't, and *why*.

[1] {skim, haleh, interran}@cs.umn.edu: Department of Computer Science and Engineering, 4-192 EE/CSci Bldg., 200 Union St. SE, Minneapolis, MN 55455, USA.

The design of the study described in this paper is motivated by our findings in previous experiments. In our first study, described in [2], we found evidence that the accuracy of surface shape perception was selectively impeded by texture pattern anisotropy in cases where the elongated texture markings were oriented in a way that was different from the first principal direction. In a followup study [3], we found that subtle surface shape differences were less easily discriminated under conditions of anisotropic texturing when the direction of the anisotropy did not follow the first principal direction. Specifically, participants were able to more reliably perceive smaller shape differences when test surfaces were textured with a pattern whose orientation followed one of the principal directions rather 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.

2. CURRENT EXPERIMENT

Having laid to rest the question of whether the orientation characteristics of an anisotropic texture pattern matter (clearly they do), in current study we try to address some of the remaining important questions about how to best define a subtle and aesthetic texture pattern that can be used to facilitate veridical shape perception without introducing unwanted visual noise. In other words, if we want to use a principal direction oriented pattern, what *kind* of principal direction texture is best, if any difference exists?

Recently we have developed a tool that allows us to map a wide class of 2D texture patterns onto arbitrary manifold surfaces, without introducing visible seams or projective distortion, and in such a way that the dominant direction in the texture pattern is constrained to follow a specified vector field over the surface at a per-pixel level [1]. While this tool allows us to apply nearly any pattern we choose, the question of what pattern to choose remains open.

In the current experiment we use performance on a surface attitude probe adjustment task to evaluate the relative effectiveness of three distinct conditions of principal direction pattern orientation in accurately conveying surface shape. The three texture conditions are: 1) a doubly-oriented texture in which a grid of approximately evenly-spaced lines follow both of the principal directions ('2-dir'); 2) a singly-oriented texture in which approximately evenly-spaced lines follow the first principal direction only ('1-dir'); and 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. For added insurance we also include a control condition of no texture, because several of the participants in our previous experiments had expressed a sentiment that "the texture seemed to be just getting in the way and making the task harder".

2.1 Objectives

The immediate goal of the current experiment is to determine whether observers are able to make more accurate surface shape judgments under some principal direction texture conditions than under others. Specifically we were interested in determining whether shape perception might be better facilitated in the condition of a texture that contains elongated elements that can be interpreted to follow both of the principal directions simultaneously than with a texture in which the elongated elements are oriented solely in one of the two principal directions. Additionally, we were interested in probing the potential effects of other texture pattern characteristics, besides orientation. In order to make the problem tractable, we chose patterns that varied in only one or two important aspects. For completeness it would have been nice to include a doubly oriented lic-like texture as well, but we had some difficulties determining how to obtain such a pattern without destroying key aspects of the perceptual equivalence of the pattern to the singly oriented sample, so we ultimately decided to leave this matter for future consideration.

2.2 Method

2.2.1 Stimulus preparation

The first step in the preparation of the experimental stimuli was to define the texture patterns that would be synthesized over the test surfaces. 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 1) from the same uniform light grey base pattern. We obtained the lic-like pattern (the rightmost texture in figure 1) by applying Photoshop's built-in motion blur filter to an input image of high frequency random noise. We then adjusted the grey levels of the three patterns to bring them as closely into overall equivalence as possible. We succeeded in achieving nearly equal luminance means (μ_{2D}=156.26, μ_{1D}=159.69, μ_{LIC}=127.04), but the line patterns still have significantly different luminance histograms and standard deviations from the lic pattern (σ_{2D}=84.37, σ_{1D}=97.93, σ_{LIC}=21.10). All three patterns span an equivalent range of spatial frequencies, but the histograms of the amplitude spectra differ in significant respects.

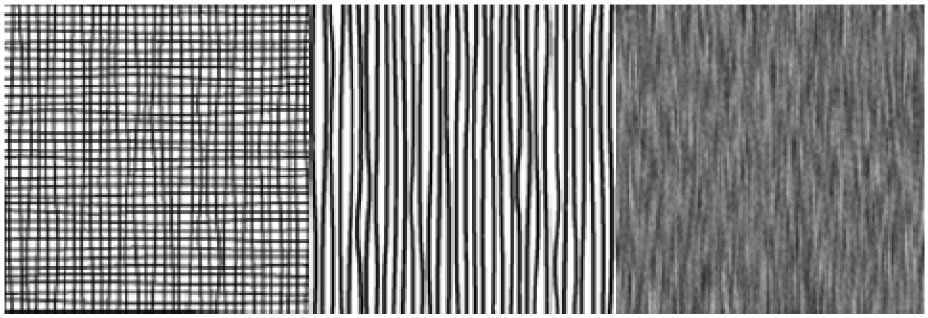

Figure 1: Sample texture patterns used in the study. From left to right: two-directional (2dir), one-directional (1dir) and lic-like (lic).

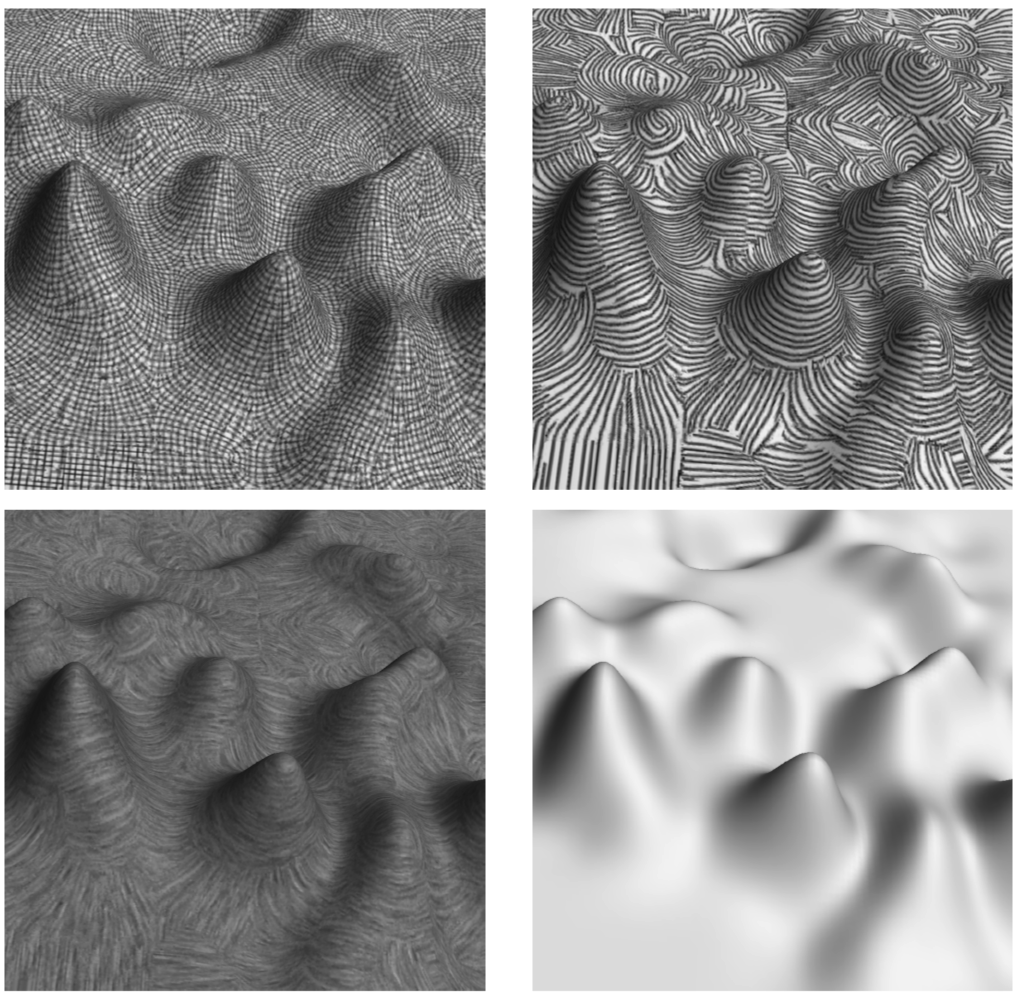

Figure 2: A test surface under each of the four texture conditions used in our experiments. Upper left: Two-directional line pattern, following the first and second principal directions. Upper right: One-directional line pattern, following the first principal direction. Lower left: One-directional LIC pattern, following the first principal direction. Lower right: No texture (control condition).

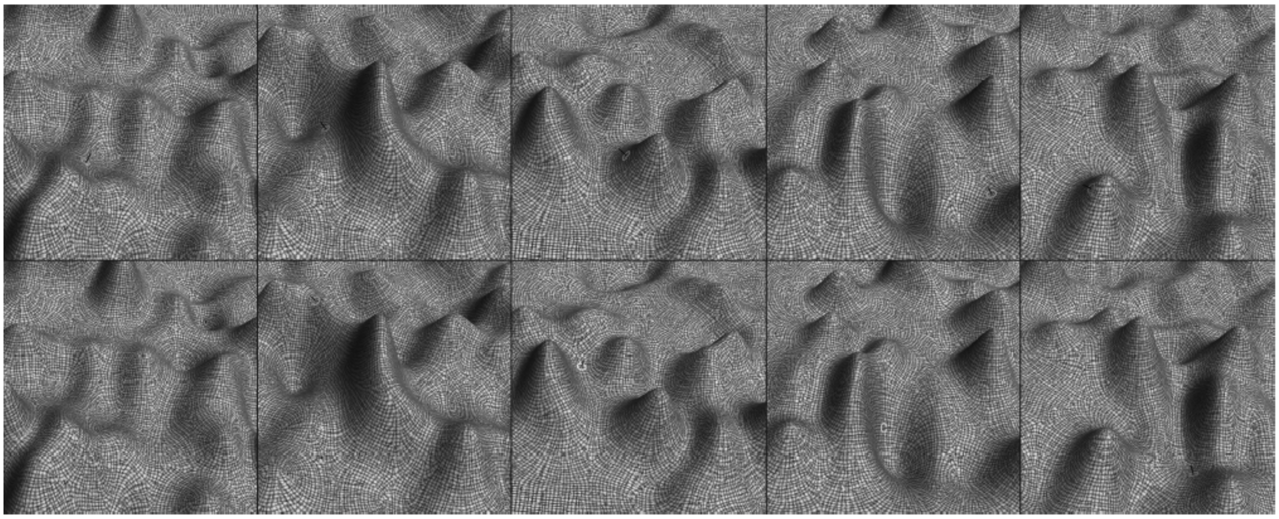

Figure 3: The five surfaces and ten probe positions used in our study. Even numbered probes 0-8 appear in the top row, with the odd numbered probes below them.

The second step in the preparation of the experimental stimuli was to define the arbitrary smoothly curving surfaces that the participants would use in making their surface shape judgments. Figure 2 shows the textures on one of these test surfaces. Following the same procedures that we used in our previous study [3], we began with a flat B-spline surface defined by a 16 x 16 grid of control points distributed at uniform intervals in the x and y directions across the z=0 plane. Over a series of 100 iterations we randomly chose a single interior control point to be perturbed by one unit (equivalent to 1/16th of the width of the surface) in either the +z or -z direction. Having the parametric definition of the B-spline surfaces, we were able to compute the first and second principal directions analytically at every vertex of the surface mesh to use in the texture synthesis.

The final step in the preparation of the stimuli was the definition of the actual surface texture, for which we used the "fitted texture" synthesis method recently developed in our lab. Since the details of this texture synthesis method are described elsewhere [1], they will be only briefly summarized here. Basically the method is a two-step process in which the surface is first split into a collection of nearly planar patches, and then the texture pattern is synthesized over each patch using the boundary conditions supplied by neighboring patches to maintain the pattern continuity across the surface. As previously mentioned, this method is capable of efficiently synthesizing unlimited quantities of a texture pattern that is perceptually equivalent to the pattern in a provided 2D sample, and does so in such a way that the resulting texture can be applied nearly seamlessly over the surface without incurring projective distortion artifacts. We synthesized each of the three test patterns over each of the five test surfaces as shown in figure 3.

2.2.2 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 in order 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.

2.2.3 Training

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.

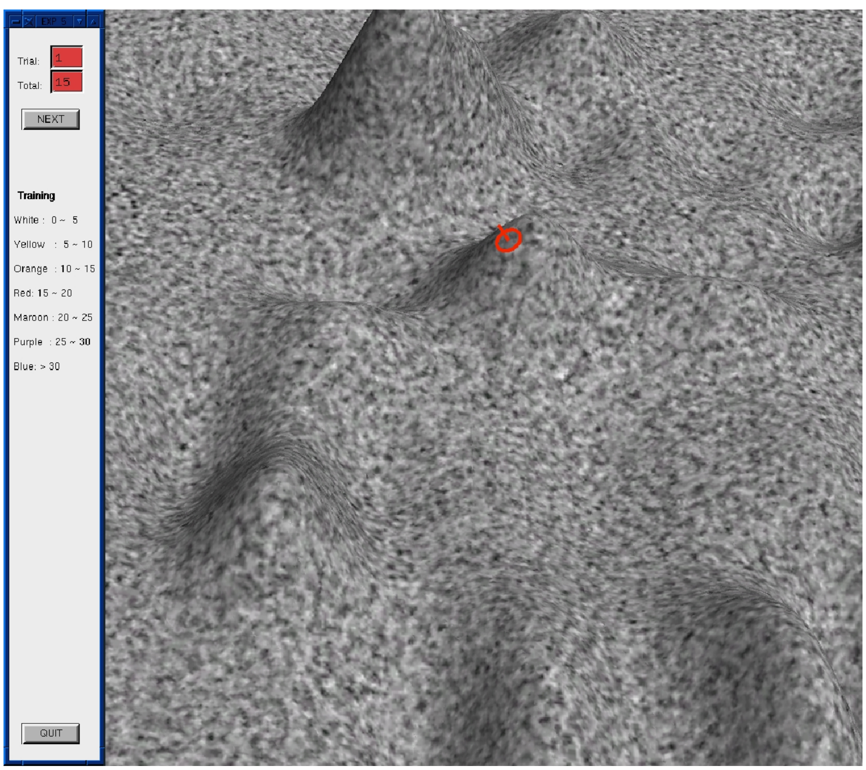

Figure 4: The training surface with one of the 15 test probes, shown in an orientation that is within 10-15 degrees of the true position.

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 4 shows an example of the training data.

2.3 Results

Figure 5 shows an overall summary of the results that we found in this experiment. Both the mean and median angle errors, taken in aggregate across all observers and all probe locations, followed this pattern. Performance was best, overall, 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', developed by Prof. Gary Oehlert from the Department of Statistics at the University of Minnesota, to perform a three-way, within subjects mixed analysis of variance (ANOVA) to evaluate the statistical significance of the results. We found significant main 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 main 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.

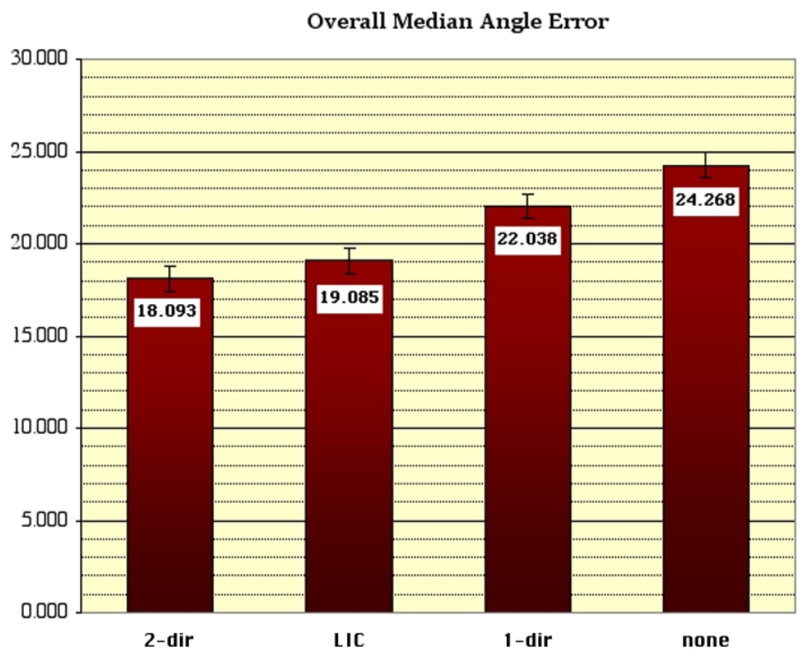

Figure 5: Median angle errors in the different texture conditions, over all subjects and all probe locations, and 95% confidence intervals.

From the charts in figure 6 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. Although remarkable consistency can be seen in the pattern of performance by texture type, across subjects, at the same probe locations, we can see from these charts that performance at individual probe locations did not always follow the same pattern as the aggregate performance shown in figure 5.

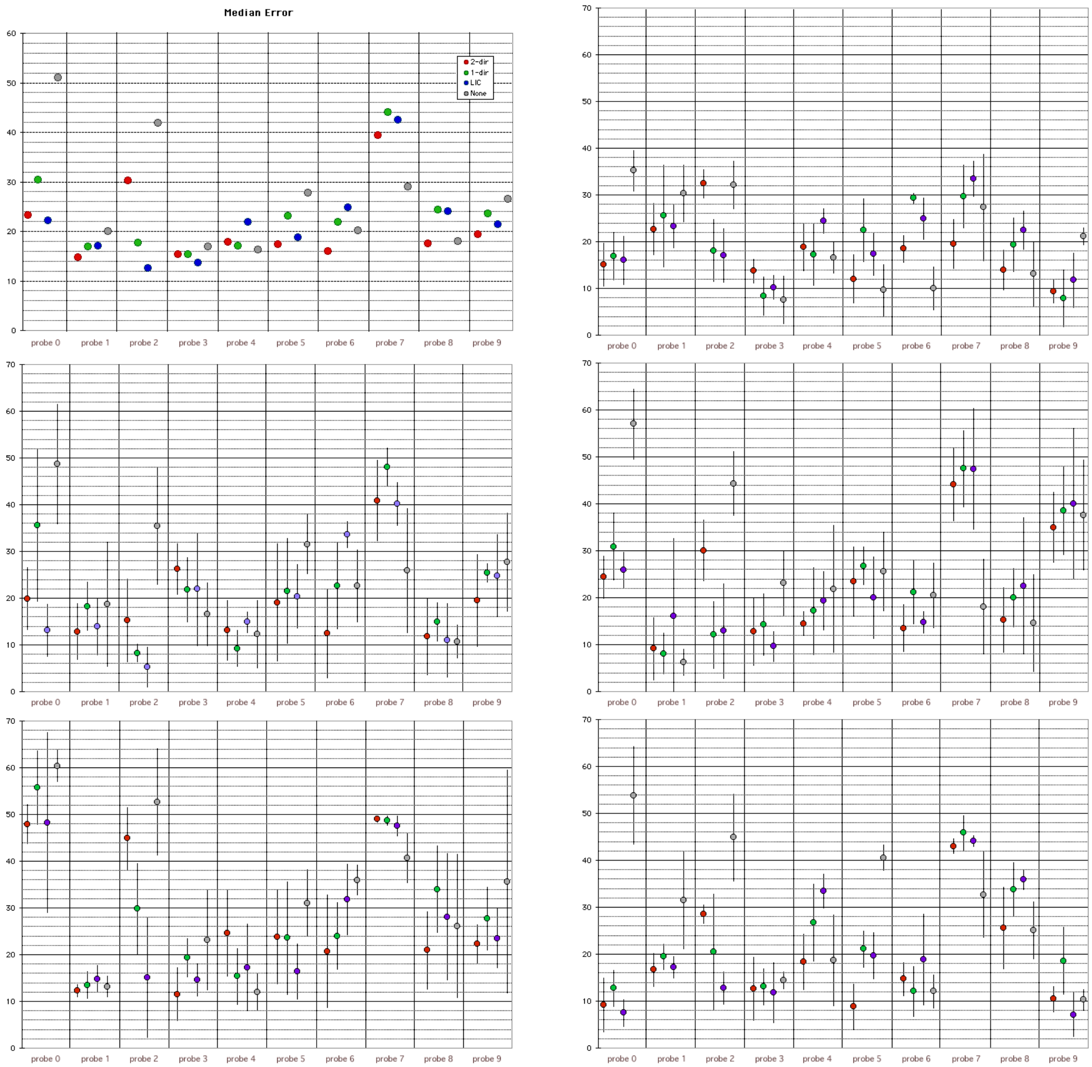

Figure 6: Charts illustrating the details of the experimental results. In the upper left, a graph of the median angle error across all subjects under each texture condition, broken down by probe location. Following that are five graphs showing the mean angle errors and standard deviations for each subject individually, again broken down by probe location. Bars show standard deviations of the results.

3. DISCUSSION

From the results of this experiment it appears 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 the patterns that carry information about distance along the principal direction than in the case of patterns that only indicate the direction itself. It could also be that performance is hindered in the case of high contrast, high regularity patterns whose statistics least resemble the statistics of patterns found in nature [4]. 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, at least in part, by one of these uncontrolled differences, rather than or in addition to the effects due to representing surface orientation information in two rather than only one direction. 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. Additionally, we suspect an interaction between pattern contrast and the accessibility of shape-from-shading information.

4. FUTURE WORK

In going through the study ourselves during the pilot phase of the experiment, we realized that the amount of global vs. local information that we were using to make our probe adjustment decisions was varying between different texture type conditions and different probe locations. We are currently conducting follow-up experiments in which we explicitly control the size of the visible area surrounding a probe, in order to investigate potential interactions between texture type and task performance as a function of this neighborhood size.

Now that we have a tool for applying any arbitrary pattern to any arbitrary surface at a high resolution while controlling the pattern orientation at a per-pixel level, we have the potential 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 easiest to control at a fine-tuned level).

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'. However in recent years, vision researchers have been discovering increasing evidence that our perceptual system is optimized for the kinds of input found in our natural environment. 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.

Also, we would like to revisit the question of determining the relative effectiveness of isotropic vs. anisotropic textures for shape representation. In earlier tests we had found no significant differences between these two conditions overall, but we suspect that these findings might have 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. By addressing the weaknesses in the principal direction texture model and incorporating some of the strengths of the isotropic texture model, it is possible that we will be able to achieve a pattern that more optimally facilitates shape perception. Investigations of the kind pursued in this study help us to determine where the most fertile ground lies for such pursuits.

ACKNOWLEGMENTS

This work was supported by a National Science Foundation Presidential Early Career Award for Scientists and Engineers (ACI-9875368). Brenda DeBlois provided us with assistance in the experimental analysis through the Statistical Consulting Clinic, which is supported by funding from the Minnesota Agricultural Experiment Station. 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. 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 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. We are also indebted to our observers CP, SL, SK, II and TU for their dedicated and conscientious efforts.

REFERENCES

[1] Gabriele Gorla, Victoria Interrante and Guillermo Sapiro (2002). "Texture Synthesis for 3D Shape Representation", *IEEE Transactions on Visualization and Computer Graphics*, to appear.

[2] Victoria Interrante and Sunghee Kim (2001). "Investigating the Effect of Texture Orientation on Shape Perception", *Human Vision and Electronic Imaging VI*, SPIE **4299**, pp. 330-339.

[3] Victoria Interrante, Sunghee Kim and Haleh-Hagh-Shenas (2002). "Conveying 3D Shape with Texture: Recent Advances and Experimental Findings", *Human Vision and Electronic Imaging VII*, SPIE **4662**, pp. 197-206.

[4] C. Alej Párraga, Tom Troscianko and David J. Tolhurst (2000). "The Human Visual System is Optimized for Processing the Spatial Information in Natural Visual Images", *Current Biology*, **10**, pp. 35-38.