Conveying 3D Shape with Texture:
Recent Advances and Experimental Findings

Victoria Interrante , Sunghee Kim and Haleh Hagh-Shenas

Department of Computer Science and Engineering, University of Minnesota

ABSTRACT

If we could design the perfect texture pattern to apply to any smooth surface in order to enable observers to more accurately perceive the surface's shape in a static monocular image taken from an arbitrary generic viewpoint under standard lighting conditions, what would the characteristics of that texture pattern be? In order to gain insight into this question, our group has developed an efficient algorithm for synthesizing a high resolution texture pattern, derived from a provided 2D sample, over an arbitrary doubly curved surface in such a way that the orientation of the 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. In this paper, we report the findings of a recent experiment in which we attempt to use this new texture synthesis method 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). In a four alternative forced choice task, we asked participants to identify the quadrant in which two B-spline surfaces, illuminated from different random directions and simultaneously and persistently displayed, differed in their shapes. We found, after all subjects had gained sufficient training in the task, that accuracy increased fairly consistently with increasing magnitude of surface shape disparity, but that the characteristics of this increase differed under the different texture orientation conditions. Subjects were able to more reliably perceive smaller shape differences when the surfaces were textured with a pattern whose orientation followed one of the principal directions 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 regardless of the surface 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, 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.

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

1. INTRODUCTION

Numerous studies have shown that shape perception can be facilitated by the presence of surface texture[13,4], but the mechanisms of texture's effect on shape perception are not yet completely understood and the question of how best to design and apply a texture pattern to a surface in order to most effectively facilitate the accurate perception of its shape remains open. 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[12,8,11,9,6,10]. The results of studies that we conducted last year[6] using a restricted class of uni-directional texture patterns[5] appeared to support the hypothesis that accurate shape perception is most severely impeded by texture anisotropy when the flow of the texture pattern turns in the surface, and that shape perception accuracy is not significantly different in the case of a unidirectional texture pattern that is locally aligned with the first principal direction than in the case of an isotropic texture pattern of similar spatial frequency. However two important questions were raised by this earlier work.

First: why, if there is little ecological justification for a texture pattern being oriented in the principal directions across a doubly curved surface, does shape perception seem to be most accurate in the principal direction orientation condition? Is it because observers are biased to interpret surface markings as being aligned with the principal directions[12,11], or is it because principal direction oriented textures intrinsically carry more shape information by virtue of their tracing out

---

Email: {interran, skim, haleh}@cs.umn.edu; Web: http://www.cs.umn.edu/{~interran, ~skim}.

lines of maximum curvature over the surface – a 3D analogy of the effect found in 2D by Biederman[1]? (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; the contour traced out by the texture flow along any other direction at that point and for the same view 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. Although these directions can be reliably classified into two types, the first principal direction and the second principal direction, it is not always clear which of these two directions a singly-oriented directional texture should follow. The first and second principal directions may "switch places" at many points, and it is not easy to reliably choose which direction to follow in order to minimize the apparent turning of the texture pattern in the surface. Might we be able to more effectively facilitate shape perception using an orthogonally bi-directional principal direction oriented pattern — one that has 90-degree rotational symmetry — rather than a uni-directional pattern that seems inevitably to exhibit artifacts at the umbilic points and parabolic lines where the first and second principal directions 'switch places'?

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 for the rendering of the surface stimuli.

1.1 Background and Previous Work

Because of the historical limitations of the capabilities of classical texture mapping software and algorithms, with few exceptions nearly all studies investigating the effect of surface texture on shape perception that have been conducted to date have been restricted either to the use of developable surfaces, which can be rolled out to lie flat on a 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 texture's effect on shape perception might be gained through studies conducted under less restrictive surface and texture pattern conditions. Hence we have been working on the development of several different types of texture synthesis and rendering algorithms to enable the creation of the stimuli essential to such investigations.

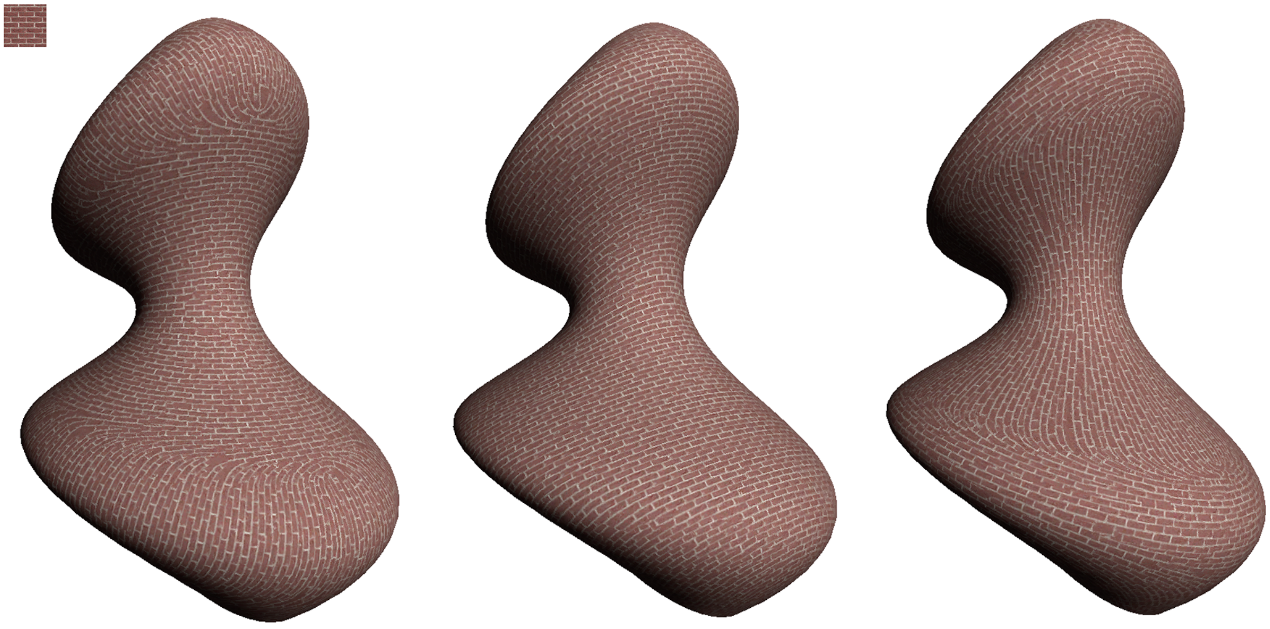

Figure 1: A brick texture pattern synthesized over a doubly curved surface according to three different texture orientation conditions. Left: rows of bricks follow the first principal direction; Center: rows of bricks are aligned in a globally uniform direction; Right: rows of bricks follow the second principal direction.

Recently, researchers in our group have developed a new algorithm capable of mapping the wide class of 2D texture patterns that can be described by Markov random processes onto arbitrary manifold surfaces, without 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. With this system we now have a means to study, in an unprecedentedly well-controlled fashion, the effects on shape perception of multiple specific texture pattern characteristics, including but not limited to orientation. Figure 1 shows some illustrative results of this new texture synthesis method, in which it is possible to informally compare 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 elsewhere[3] and 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 between the adjacent patches. Figure 2 shows several intermediate results of the texture synthesis process for one of the sample surfaces used in the experiment described in this paper.

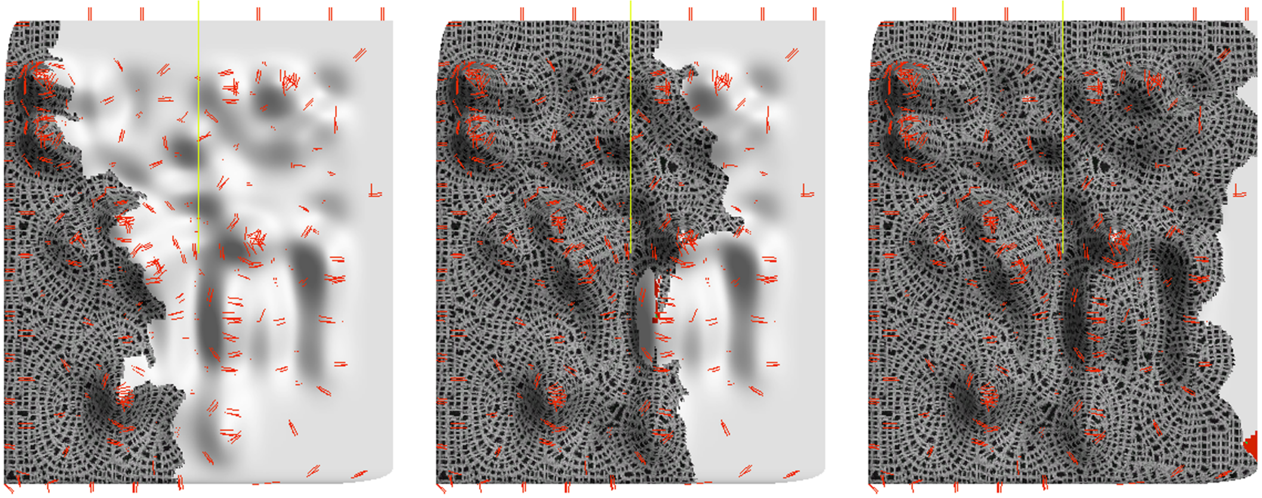

Figure 2: The process of texture synthesis illustrated on a sample surface (the red lines indicate the direction of the vector field at selected vertices).

2. METHODS

The two objectives of the experiment that we describe in this paper were: 1) to evaluate the validity of the hypothesis that shape perception accuracy declines in non-principal direction texture orientation conditions because less shape information is available to observers in those situations, and 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 which of the two principal directions to align the texture pattern with at each point. We decided to use a four alternative forced choice surface shape discrimination task primarily for the first purpose, hypothesizing that if it is true that in the case of non-principal direction oriented patterns there is less information available to be used to make shape judgments, then we should find a higher threshold for the detection of changes in surface shape under these conditions due to the greater resulting ambiguity between multiple surfaces of different shapes which might all appear to be similar. In the following sections we provide the details of the experimental set up and design.

2.1 Stimuli

Beginning with 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 in which 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, alternately. 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 be changing 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 bumps, 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 of difference. 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 3 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.

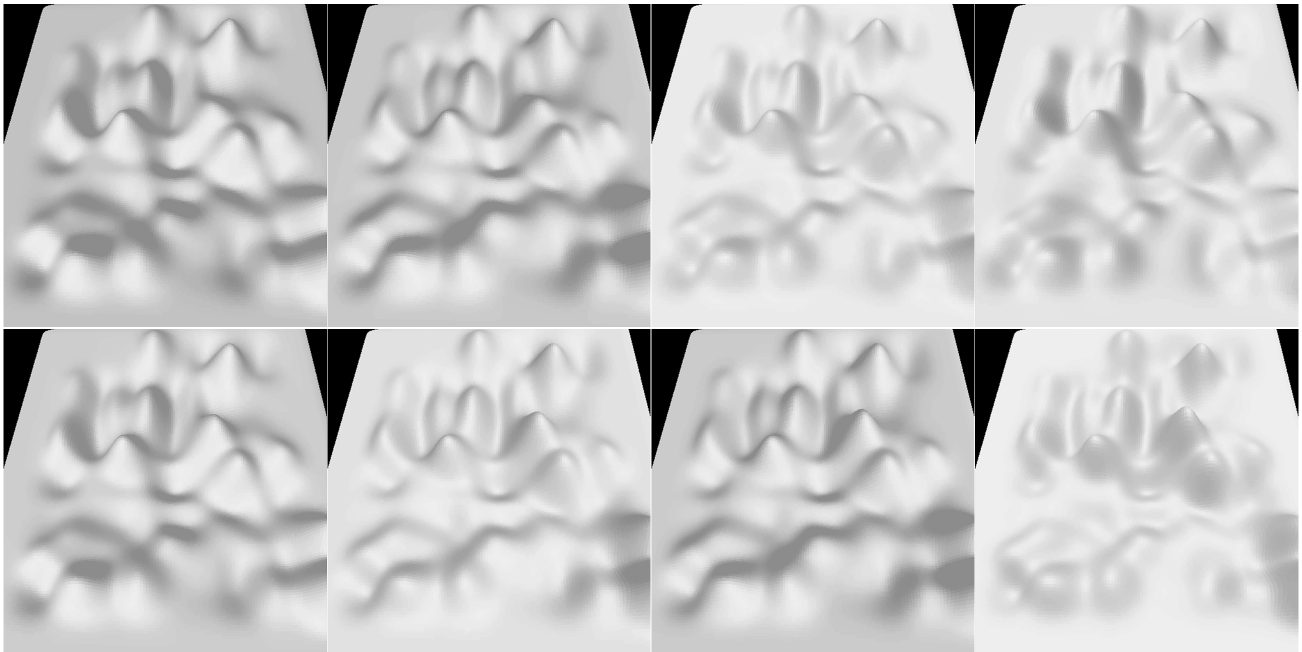

Figure 3: Shown from upper left to lower right are the eight different displacement surfaces used to represent shape changes in quadrant 1 (upper right portion of the surface). We used a different randomly selected illumination direction in each case in order to encourage our participants to focus on the 3D shape information and to minimize the usefulness of simple 2D pixel difference cues.

Once we had defined the 32 sample surfaces (4 quadrants x 8 levels of displacement on the selected feature in each), the next task was to define the three different vector fields that would control the orientation of the texture patterns over each surface. The three different texture orientation conditions that we wished to compare were: 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; the uniform direction condition, in which the texture was constrained to remain oriented in a constant uniform direction across the surface, and the swirly direction condition, in which the orientation of the texture pattern was allowed to twist and turn in a sinusoidal manner over the surface.

Having the parametric definition of the Bspline surfaces, we were able to compute the first and second principal directions analytically at every vertex in the surface 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.

The final step in the preparation of the stimuli was the surface texturing, for which we would use the new 'fitted texture' synthesis method developed in our lab. 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 doing so in such a way that the resulting texture can be applied nearly seamlessly over the surface without incurring projective distortion artifacts. We began with two base texture patterns selected from the Brodatz texture album[2]. As the first texture we chose a simple pattern, D49: Straw Screening, that was as plainly unidirectional as possible, and as the second texture we selected the pattern that seemed to use to be as similar looking to the first as possible while also being bi-directional, D20: French Canvas. We divided each of these patterns into eight subimages, in order to have the surface texture created from a different but similar subsample pattern for each of the different surface displacement intervals. It was necessary to take this precaution in order to avoid unwanted pixel-by-pixel texture similarities in areas of the surface that did not undergo a shape change. Figure 4 shows the same series of surfaces presented in figure 3, with the bi-directional texture pattern applied following the orthogonal principal directions. Figure 5 shows the results of using the three different texture orientation conditions (uniform, principal and swirly) over the same underlying surface.

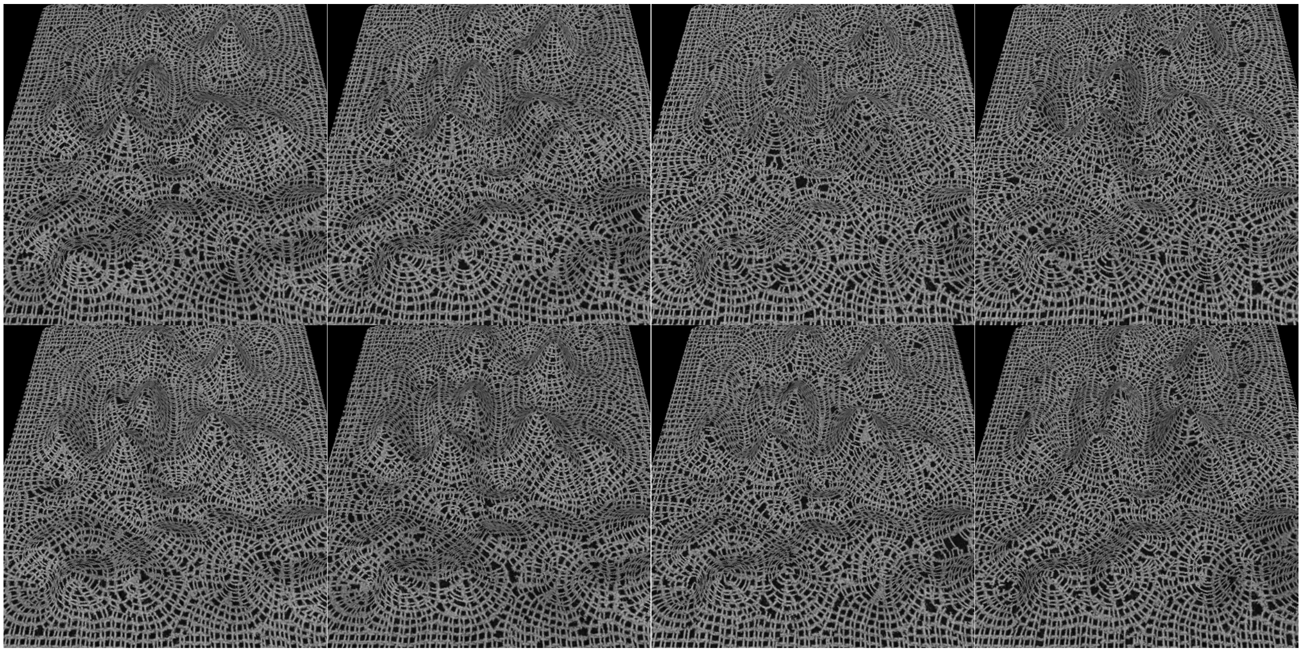

Figure 4: From upper left to lower right, the same eight different displacement surfaces, textured with principal direction oriented patterns derived from eight different subsample swatches from the same larger original texture pattern image.

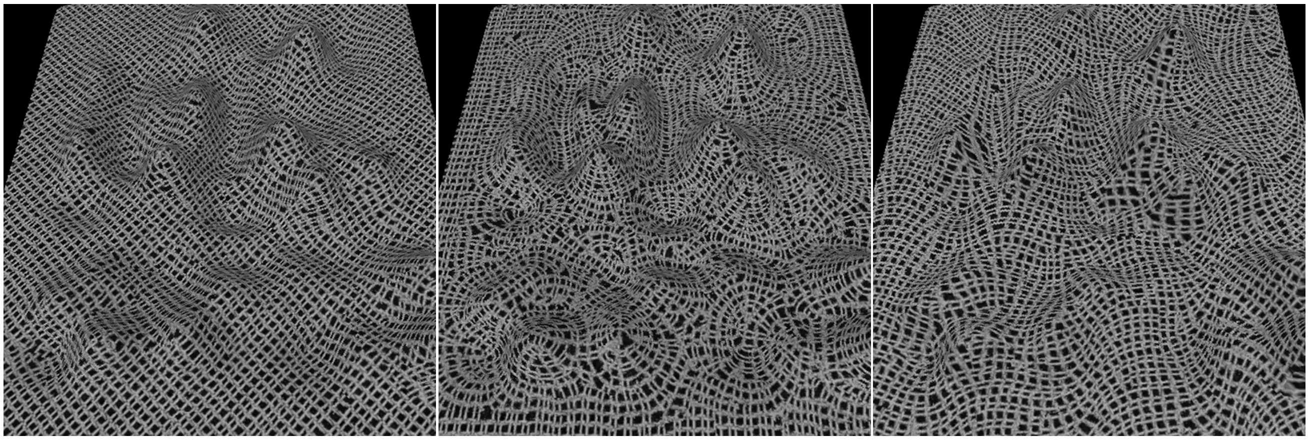

Figure 5: The three different texture orientation conditions. Left: uniform direction; Center: principal direction, Right: swirly direction.

Because we thought it was important to try to avoid the possibility of any unfortunate biases due to the choice of viewing angle, we chose to use 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.

2.2 Task

Over the course 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 6 and 7 show the user interface for the experiment with representative examples of the remaining conditions. When the user clicks on the "toggle lines" button, the boundaries between the quadrants would be explicitly drawn over the image (the boundaries were defined in 3D). At the beginning of the experiment, the lines were turned on by default. Once the user elected to turn the lines off they would remaine off for subsequent trials unless the toggle lines button was pressed a second time to turn them on again. We have elected to show the interface in the lines-off condition here in order to simplify the presentation.

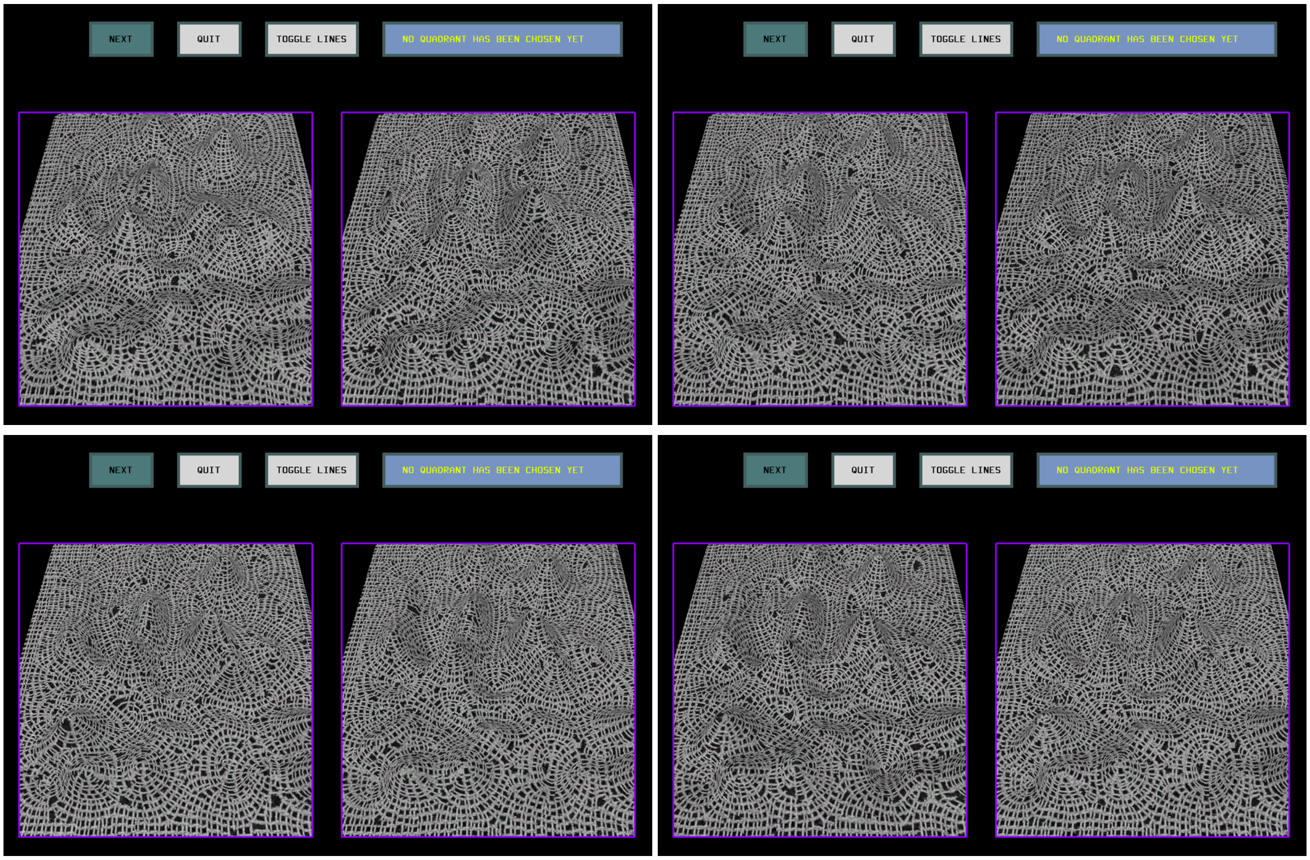

Figure 6: Four screen shots from the experiment, showing the four different shape difference types. The bi-directional pattern and the principal direction orientation condition were used in all of the cases shown here. Clockwise from the upper left, the surface shape differences appear in the following quadrants of each image respectively: upper right, upper left, lower right and lower left quadrant.

Because of the apparent simplicity of the task, we did not anticipate the need to have our subjects go through a training phase before beginning the experiment. In retrospect we discovered that this was a mistake, as we will describe in the following section.

We collected data from three subjects, who were chosen because of their availability and known reliability. One of the subjects was an author of this paper, who was intimately familiar with both the task and the experimental objectives but who was for obvious reasons not involved in the preparation of the sample surfaces used. The other two subjects were professionals from outside of our lab and outside of computer science, who were new to the task but informally aware of our basic research objectives.

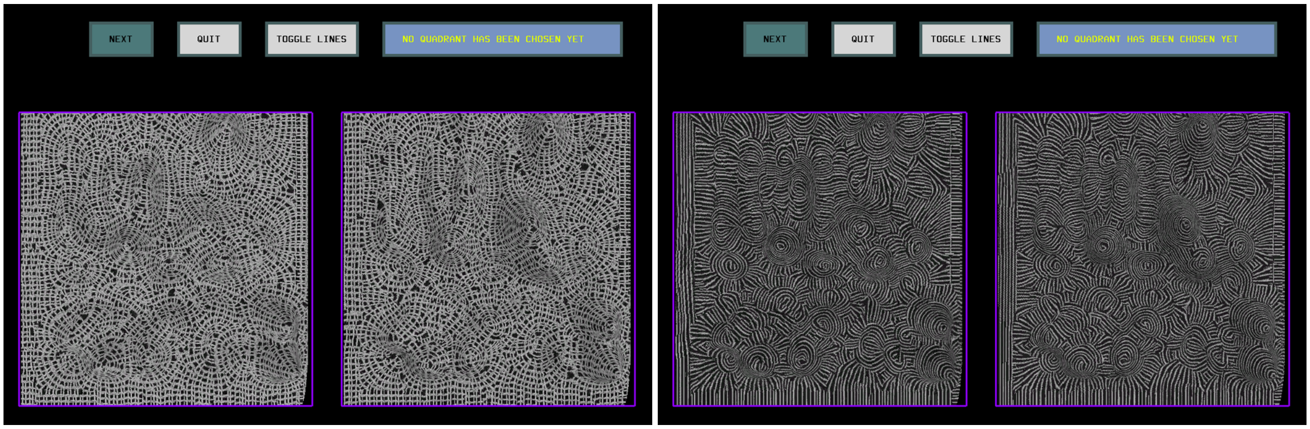

Figure 7: Examples of the two remaining conditions: the frontal view, and, on the right, the unidirectional texture pattern in the first principal direction orientation condition.

3. RESULTS

Figure 8 summarizes the overall results of our pilot 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, and to make these graphs we combined the data across the two viewing conditions, the two texture pattern conditions, and the four quadrant/shape types.) We did not find it illuminating to consider each of the texture pattern conditions separately because the results turned out to be so similar in both cases. The pattern of responses seems similar across the three subjects, though the level of performance seems higher, overall, for the more experienced subject. Accuracy rates seem to rise as the disparity of the compared surface 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.

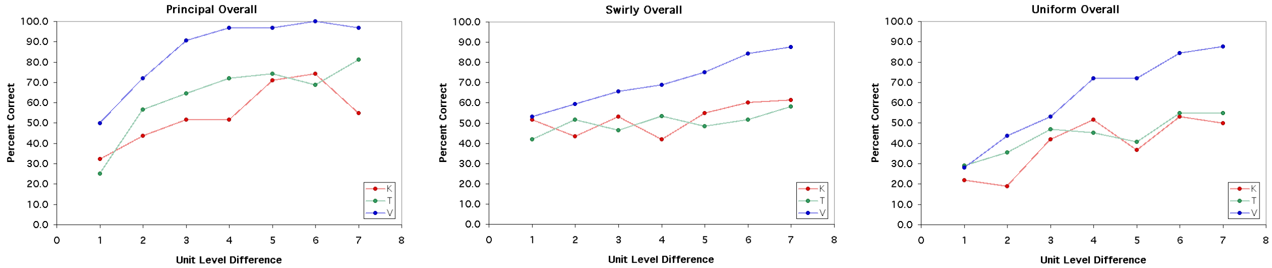
Principal Overall

Swirly Overall

Uniform Overall

Figure 8: Overall results from our pilot experiment, per texture orientation condition. Each line represents a different subject.

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 look separately at the first and second half of the data collected. This is of course not an ideal way to break down the analysis, but it proved to be an illuminating exercise. Figures 9-10 show the results tabulated separately for the first 336 responses and the second 336 responses per subject.

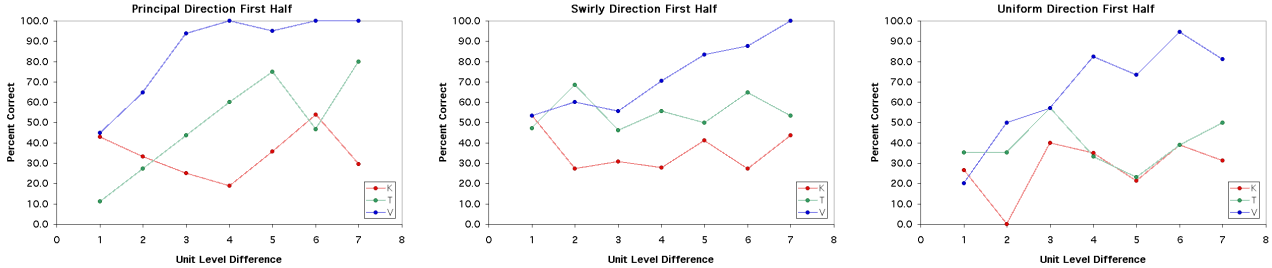
Principal Direction First Half

Swirly Direction First Half

Uniform Direction First Half

Figure 9: Results from the ‘first half’ only, per texture orientation condition.

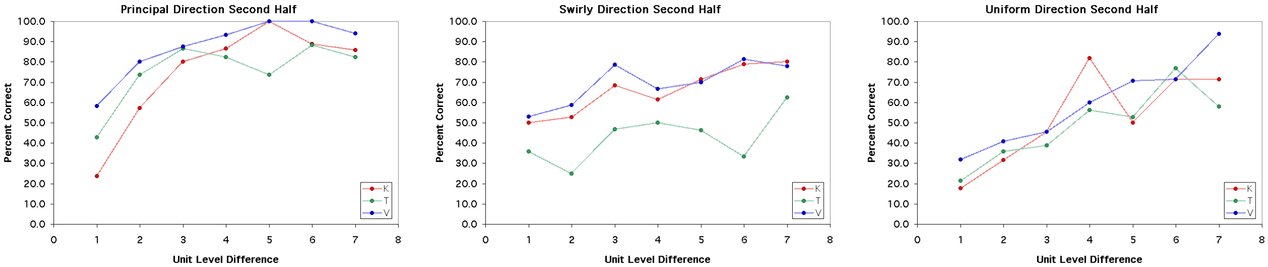

Figure 10: Results from the ‘second half’ only, per texture orientation condition.

It is clear from the data in figures 9 and 10 that something very different is happening in the first and second halves of the experiment, at least in the case of the two observers less familiar with the task and the subject. Noting that the three observers’ performance converges well in the second half, but not well at all in the first, we suspect that some considerable unreliability is being introduced into the results at the beginning of the experiment, most likely due to confusion on the part of the observers about how to interpret all of the different kinds of differences that appear between the two images whose shapes are supposed to be compared. Because only four different distinct shape features are actually changing, it’s likely that all of the observers eventually figured out how to differentiate between the apparent image differences that were due to the vagaries of the texture synthesis process, which was working from a different sample image in each case, and the differences that were 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 11, 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 experienced participant, 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, and 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. Unfortunately, because so much varied information has been combined to obtain these numbers, it is difficult for us to be confident of the significance of these apparent differences.

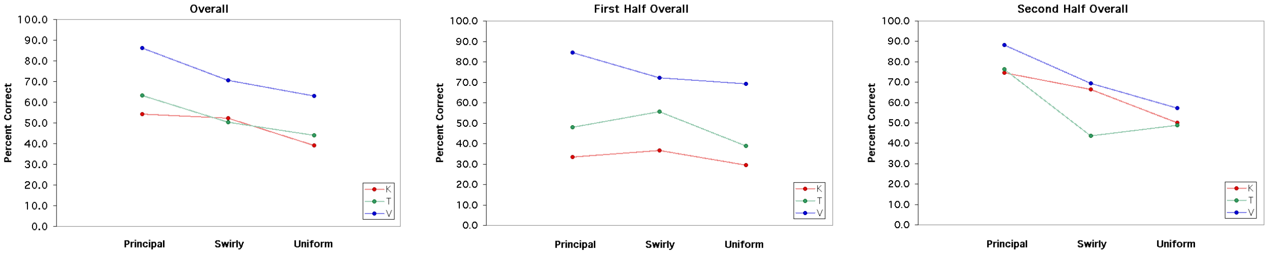

Figure 11: Combined results (averaged over all displacement ranges), per texture orientation condition.

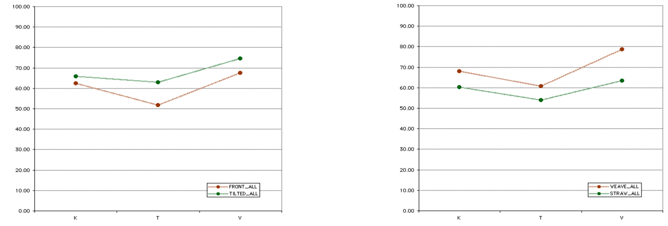

Figure 12: A comparison of performance under the two viewing conditions (left) and the two texture type conditions (right).

Figure 12 summarizes the overall differences, or lack thereof, in the cases of the two viewing conditions and the two texture type 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 not by much.

4. DISCUSSION

The results the we obtained in this experiment seem to lend support to the hypothesis that principal direction oriented texture patterns have a potential advantage over non principal direction oriented pattern in facilitating shape perception because of the fact that they provide more explicit evidence of the potential amount of surface curvature present than do directional patterns oriented over the surface 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 than explicitly followed both the first and second principal directions simultaneously rather than following only one of these directions. We speak more about this in the future work section.

One question that is raised by this experiment is, why did we seem 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 in-surface 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 discriminabilty 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. Because 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 have earlier noted, 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.

4.1 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 in some extreme cases be partially hidden from view of the observer due to occlusion by other parts of the surface. As it turns out, in the 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 randomly vary the illumination direction 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, shape perception has been shown to be facilitated to a greater extent when the incident light strikes the surface at a grazing angle than when the light hits the surface head-on[7]. 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 to have the only difference in these three cases be the mode of orientation of texture over the surface, with the random selection of lighting direction and random selection of difference interval endpoints being the same for all three texture conditions in that particular one unit shape difference case. Unfortunately that is not how it turned out that things were managed in this experiment. Too late, 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 have every reason to 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 followup experiment in which these extraneous random variations are more carefully controlled for.

5. FUTURE WORK

While it is logical to assume that observers' ability to accurately perceive surface shape is poor when they are unable to accurately identify 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 accurately identify the quadrant in which two presented surfaces differ in shape, that this is because they are in fact able to accurately perceive the two shapes. 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 are not shaped the same. Hence, the conclusions that can be 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, because they are defined according to surface shape properties, would exhibit more prominent variation in response to shape changes than the other two types of textures, whose relationship to 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.

6. ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation (NSF ACI-9875368). We are grateful to Paul Schrater and Cindee Madison for many helpful suggestions on the experimental design. The texture synthesis and image rendering was done using software originally written by Gabriele Gorla. Prototype surface modeling code was written by Krista Janssen, with support from the CRA-W Distributed Mentor Project. We are grateful to our observers TF and KB for their dedicated efforts.

7. REFERENCES

1. I. Biederman. "Human Image Understanding: Recent Research and a Theory", in Human and Machine Vision, Azriel Rosenfeld, ed., Academic Press, pp. 13-57, 1985.
2. P. Brodatz. *Textures: A Photographic Album for Artists and Designers*, Dover Publications, 1966.
3. G. Gorla, V. Interrante and G. Sapiro. "Texture Synthesis for 3D Shape Representation", *IEEE Transactions on Visualization and Computer Graphics, to appear.*
4. V. Interrante, H. Fuchs and S. Pizer. "Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture", *IEEE Computer Graphics and Applications*, **3**(2): 98–117, 1997.
5. V. Interrante. "Illustrating Smooth Surfaces in Volume Data via Principal Direction-Driven 3D Line Integral Convolution", *ACM SIGGRAPH Proceedings*, pp. 109–116, 1997
6. V. Interrante and S. Kim. "Investigating the Effect of Texture Orientation on Shape Perception", *Human Vision and Electronic Imaging VI*, SPIE **4299**, pp.. 330-339, 2001.
7. A. Johnston and P. J. Passmore. "Shape from Shading I: Surface Curvature and Orientation", *Perception*, **23**, pp. 169-189, 1994.
8. D. C. Knill. "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.
9. A. Li and Q. Zaidi. "Perception of Three-Dimensional Shape from Texture is Based on Patterns of Oriented Energy", *Vision Research*, **40**, pp. 217–242, 2000.
10. A. Li and Q. Zaidi. "Information Limitations in Perception of Shape from Texture", *Vision Research*, **41**, pp. 1519-1534, 2001.
11. P. Mamassian and M. P. Landy. "Observer Biases in the 3D Interpretation of Line Drawings", *Vision Research*, **38**, pp. 2817–2832, 1998.
12. K. Stevens. "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.
13. J. T. Todd, J. F. Norman, J. J. Koenderink and A. M. L. Kappers. "Effects of Texture, Illumination and Surface Reflectance on Stereoscopic Shape Perception", *Perception*, **26**(7), pp. 807-822, 1997.