Conveying Shape with Texture: an experimental investigation of the impact of texture type on shape categorization judgments

Sunghee Kim, Haleh Hagh-Shenas and Victoria Interrante

University of Minnesota

Abstract

As visualization researchers, we are interested in gaining a better understanding of how to effectively use texture to facilitate shape perception. If we could design the ideal texture pattern to apply to an arbitrary smoothly curving surface in order to enable its 3D shape to be most accurately and effectively perceived, what would the characteristics of that texture pattern be? In this paper we describe the results of a comprehensive controlled observer experiment intended to yield insight into that question.

In previous studies, using a surface attitude probe adjustment task, we had found that judgments of shape under conditions of anisotropic texturing were most accurate when the directions(s) of anisotropy were aligned with one or both of the principal directions of curvature over the surface, as opposed to being aligned with an arbitrary constant direction over the surface, or with a direction that varied over the surface in a manner unrelated to the surface geometry. However, many open questions remained.

Here we report the results of a new study comparing the relative accuracy of observers' judgments of shape type (elliptical, cylindrical, hyperbolic or flat) and shape orientation (convex, concave, both, or neither) for local views of boundary masked quadric surface patches under six different principal direction texture pattern conditions plus two control texture conditions (an isotropic pattern and a non-principal direction oriented anisotropic pattern), under both perspective and orthographic projection conditions and from both head-on and oblique viewpoints.

Our results confirm the hypothesis that accurate shape perception is facilitated to a statistically significantly greater extent by some principal direction texture patterns than by others. Specifically, we found that, for both views, under conditions of perspective projection, participants more often correctly identified the shape category and the shape orientation when the surface was textured with the pattern that contained oriented energy along both the first and second principal directions only than in the case of any other texture condition. Patterns containing markings following only one of the principal directions, or markings oriented obliquely to the principal directions, or containing information along other directions in addition to the principal directions yielded poorer performance overall.

In examining the effects of projection type and view direction, we found that observers retained the ability to make correct shape category judgments under conditions of orthographic viewing under many texture conditions, as long as the surface patches were viewed from an oblique vantage point rather than head-on. However, our observers were unable to reliably disambiguate convex from concave surface orientations in the absence of perspective projection, regardless of texture type.

CR Categories and Subject Descriptors: I.3.7, I.2.10, J.4, H.5.2.

Additional Keywords: shape perception, texture, principal directions

1. INTRODUCTION

As visualization designers, our goal is to determine how to most effectively portray a set of data such that its essential features can be easily and accurately understood. When we use computer graphics techniques to display computed or acquired surfaces, we have wide discretion over the choice of the surface material properties. If we desire to portray a surface in a way that best facilitates the accurate, intuitive understanding of its 3D shape, what rendering characteristics should we choose to most effectively accomplish this task? The answer to this question has significant potential relevance to a wide range of visualization applications in which scientists need to attain an accurate, intuitive understanding of the shapes of complicated, smoothly curving surfaces in their data. The most common practice in rendering objects is to use a simple Phong shading model without any surface texture. Phong shading is frequently used because it is easy to implement and is the default on most systems. However, smooth shading is not optimal for all purposes and in particular is not optimal for shape representation — research in shape perception has shown that shape understanding can be facilitated by the presence of the right kinds of surface texture. Unfortunately, existing theories do not provide sufficient guidance to definitively answer the question of how best to define a surface texture pattern to meet this goal. Over the past several years, we have conducted a series of experiments [4, 5, 6] investigating the impacts of various characteristics of surface texture patterns on shape perception. In this paper we present the findings of our most ambitious and successful experiment to date. But first we provide a brief overview of previous work in shape perception from shading and texture.

2. PREVIOUS WORK

Observation tells us that shading clearly plays an important role in conveying information about the shape structure of a surface. However, psychophysical experiments have indicated striking limitations in observers' ability to accurately infer some types of shape information solely from the pattern of diffuse shading over a local, smoothly curving surface patch due to illumination by a single light source. In a study on the perception of local surface orientation from shading [11], Mamassian and Kersten presented observers with images of simple smooth objects under four illumination conditions plus a silhouette control condition

---

4-192 EECS Bldg, 200 Union St. SE, Minneapolis, MN 55455
{skim, haleh, interran}@cs.umn.edu

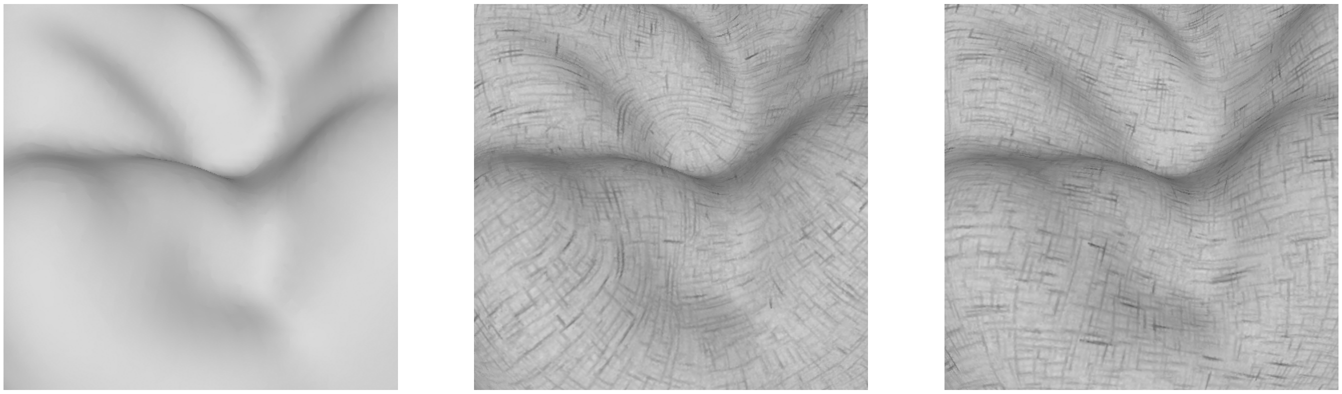

Figure 1: A close-up view of the top portion of a tooth dataset, depicted, from left to right, with: no texture, an orthogonal texture pattern following smoothed principal directions, and an orthogonal texture pattern following constant uniform directions in object space.

(uniform ambient shading) and asked them to make judgments about the local surface orientation at selected points. Analyzing the results using an ideal observer analysis, they concluded that the human observers appeared not to be using the shading information at all, but rather relying on the occluding contours of the objects as the overriding basis for their judgments. In an experiment conducted by Erens and colleagues [1], in which participants were shown boundary-masked, computer-generated images of quadric surface patches under a variety of lighting conditions, observers were unable to correctly infer the direction of illumination, or to accurately categorize the shape of the patch as either elliptic or hyperbolic. When, in a follow-up study, the direction of illumination was explicitly indicated, observers gained the ability to disambiguate convex from concave surface orientations, in the elliptic case, but were still unable to reliably differentiate the shape type. Thus, Erens *et al* concluded that local shading structure, by itself, is only a weak cue to surface shape.

In this research, we set out to answer the questions: Is it possible that observers will be able to reliably discriminate between elliptic, hyperbolic, and cylindrical patches under conditions where surface texture – in the form of a pattern of luminance variations - is present in addition to shading? Will it be the case that shape category identification is enabled under some texture conditions but not others? If this were to be the case, it would provide us with a useful method for differentiating texture patterns that have a greater potential to be helpful in facilitating shape perception from texture patterns that do not.

Despite decades of research, investigations into the effects of texture on shape perception continue to be of interest and importance in the vision research community to this day. Because reliable computer graphic techniques for applying an arbitrary given texture pattern to an arbitrary doubly curved surface have only recently been developed [e.g. 2], research on shape perception from texture has generally been somewhat restricted, either to developable surfaces [7, 9] (surfaces which can be rolled out to lie flat on a plane), to patterns projected onto surfaces from a particular direction [17], or to solid textures [16] (whose particular characteristics are generally independent of the shape of surface carved out of them). There are many open questions that still exist about the impact of surface texture on shape perception. However, many important insights have also been achieved.

Numerous studies have found evidence that the accuracy of observers' judgments of surface orientation and curvature can be significantly affected [both positively and negatively] by the presence of a surface texture pattern [3, 16]. Recent findings support the idea that the facilitating effects of the presence of texture depend not only upon the intrinsic characteristics of the texture pattern itself [e.g. 13] but also upon how the pattern is laid down over the surface [4, 9, 10].

Several researchers have suggested that observers may be biased toward interpreting lines on surfaces as if they were following the principal directions [12, 14]. Li and Zaidi [9, 10] have argued that two conditions are necessary for the perception of 3D shape from texture: 1) the texture pattern must have a considerable amount of energy along the direction of maximum curvature and 2) the surfaces must be viewed with noticeable perspective. However the task that they used to judge shape perception, discriminating which of two adjacent points is more distant, actually only provides coarse information about the perceived direction of surface slant, which is useful for determining whether observers can differentiate convexities from concavities, but does not capture all of the information that we would like to know about shape perception.

Other researchers [8, 15] downplay the importance to shape understanding of specific texture pattern characteristics such as alignment with the principal directions, arguing that these conditions are not always the necessary factors in conveying information to observers and demonstrating that surface shape can be reliably inferred from a very wide range of texture patterns. Appearing to contradict Li and Zaidi [10], Todd and Oomes [15] show that there is some shape information available under orthographic projection. They also describe examples in which texture elements appear able to reveal the underlying shape of an object even though the texture itself lacks significant energy in any particular direction. They argue that surfaces which do not have gradual orientation changes relative to the viewing direction are degenerate for providing information about 3D shape from gradients of texture compression.

A complicating factor in this debate is the lack of standard reliable universally accepted metrics for evaluating shape perception. Various tasks that have been used in the past are: manipulation of a surface attitude probe, indicating an estimate of the direction in which the surface normal is pointing, individually measured at a single location on the surface [17], determination of which of two points is farthest away, qualitatively indicating whether a surface appears to be tipping forward or backward in the direction between the two points [9], identification of the quadrant in which two surfaces differ in shape [5], and identification of the shape category of a surface patch [1].

In our own previous work we have found indications that surface attitude judgments are significantly more accurate [4] and surface shape discrimination thresholds significantly reduced [5] under conditions of principal direction texturing, as compared to conditions of texturing with an anisotropic pattern whose orientation is either uniform in object space or follows a non-geodesic path within the surface; in this study we sought to compare alternative principal direction oriented patterns. In particular, we sought to determine whether a pattern containing oriented elements aligned with both the first and second principal directions would show shape more effectively than a singly-oriented anisotropic pattern aligned only with the first or the second principal direction. In a previous study [6] we had found indications that this might be the case, but our results were below significance at the 95% level. In the current experiment we also sought to investigate the impact on texture perception of employing patterns containing elements systematically oriented at an oblique angle (45 degrees) to the principal direction(s). Such textures would contain information that implicitly encodes the principal directions with as great a reliability as 'unrotated' principal direction textures, but the eye would be drawn by these texture to follow lines over the surface that were oriented in a direction not equal to the principal directions, and the possibility would exist that observers might interpret the oblique lines as if they were oriented in the principal directions even though they were not. Of secondary importance but also of some interest to us was determining whether it was in fact true that accurate shape perception could not be achieved except under conditions of perspective projection, as claimed by Li and Zaidi. Like Todd, we had accumulated some anecdotal evidence that shape perception might be still possible even under conditions of orthographic viewing, and we were interested to pursue this question a little further.

3. EXPERIMENT

3.1 Method

3.1.1 Stimuli Preparation

We used eight patterns to texture our surfaces using the texture synthesis method developed by Gorla and colleagues [2]. The eight patterns used were as follows: 1-directional, (which shows the first principal direction), 1-directional rotated clock-wise 45 degrees (diagonal to the first principal direction), 1-directional rotated clock-wise 90 degrees (which shows the second principal direction), 2-directional (which shows both first and second principal directions), 2-directional rotated clock-wise 45 degrees, 3-directional (showing the first and second principal directions plus a diagonal direction), swirly and noise. Figure 2 shows example swatches of each of these patterns.

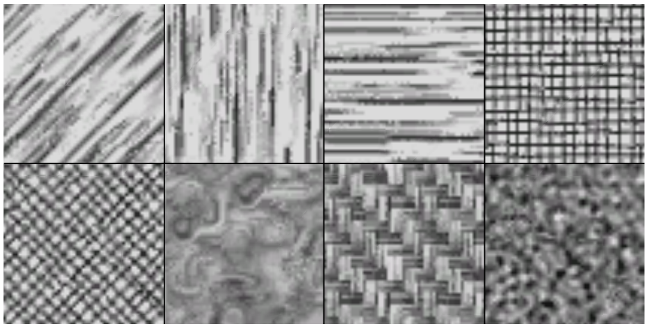

Figure 2: The sample texture patterns used in the study. Top row, from left to right: 1dir45, 1dir, 1dir90, 2dir. Bottom row, from left to right: 2dir45, swirly, 3dir, noise.

Our surface stimuli were simple height fields that we built by applying the parametric equations of each surface category to an equilateral mesh. These parametric equations represented the following surface types: ellipsoid, elliptical cylinder, saddle, and flat. The principal directions at each point of these meshes were calculated using the u and v parameters that were used in the parametric formulas, and both sides of the surfaces were textured.

3.1.2 Experimental Setup and Task Description

The study consisted of 592 trials. Observers classified each image as belonging to one of the four shape categories: ellipsoid, cylinder, saddle, or flat, and one of four shape orientations: convex, concave, both (as in the saddle), or neither (as in the flat case). We used two viewing conditions: straight-on and oblique, and two projection conditions: perspective and orthographic. To avoid the potential of uncovering orientation dependent effects we also rotated our pictures in the image plane over repeated trials, using two rotations – 0 and 90 degrees – in the case of the symmetric straight-view images, and four rotations – 0, 90, 180 and 270 degrees – in the case of the oblique views. Figures 3-11 show examples of the surface stimuli. Note that the oblique viewing direction used in the convex cases is not the direct inverse of the oblique viewing direction used in the concave cases. Although this arrangement is less than ideal, making this allowance gave us greater flexibility to satisfy the concurrent constraints that the contour of the surface patch not be visible in the image, and that the total amount of surface curvature across the visible portion of the patch be as large as possible.

The experiment was conducted in a laboratory located on the campus of University of Minnesota. During the experiment the images were shown on a 21-inch CRT monitor, one at a time. The pixel resolution of the monitor was 1600x1200. Image resolution was 1000x1000. Observers freely viewed the images under standard room lighting conditions. There was no time limit associated with the trials.

A total of 8 observers – 5 males and 3 females, ranging in age from 17-50, – participated in the study. Five were naive to the purposes of the experiment, and were compensated for their conscientious efforts. Among this group was a high school student and a professional graphic artist. The other three were members of our research team. (In analysis we found no significant differences in patterns of performance between

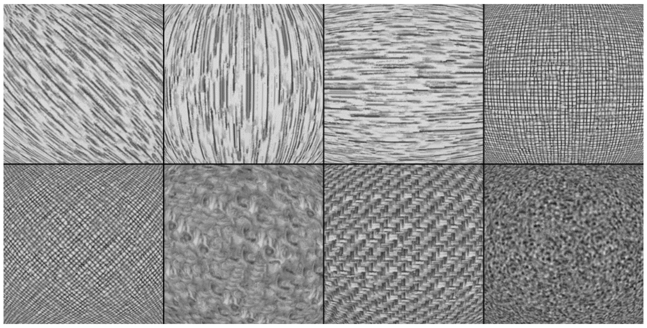

Figure 3: Convex ellipsoids in perspective projection viewed head-on, with each texture type applied.

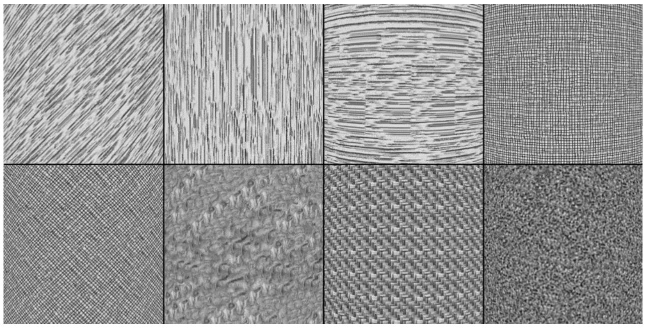

Figure 4: Convex cylinders in perspective projection viewed head-on, with each texture type applied.

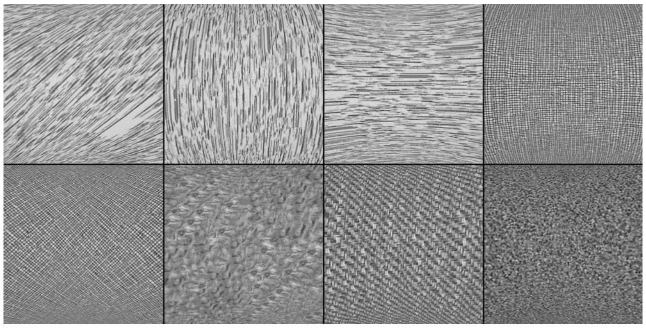

Figure 5: Saddle surfaces in perspective projection viewed head-on, with each texture type applied.

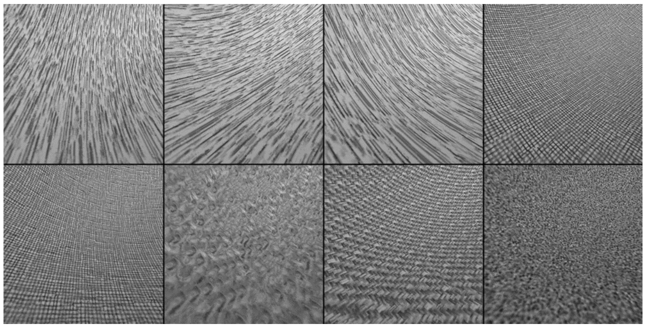

Figure 6: Concave ellipsoids in perspective projection viewed obliquely, with each texture type applied.

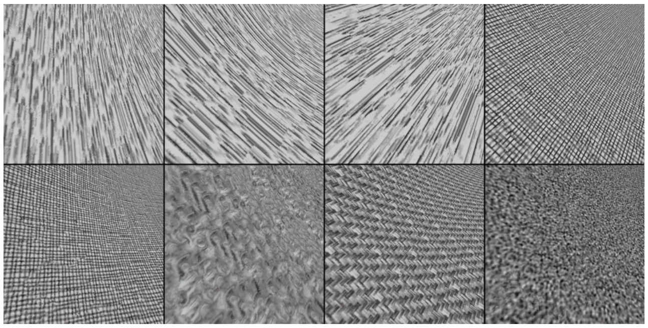

Figure 7: Concave cylinders in perspective projection viewed obliquely, with each texture type applied.

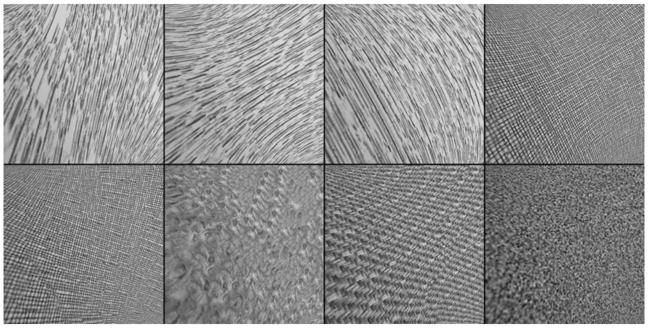

Figure 8: Saddle surfaces in perspective projection viewed obliquely, with each texture type applied.

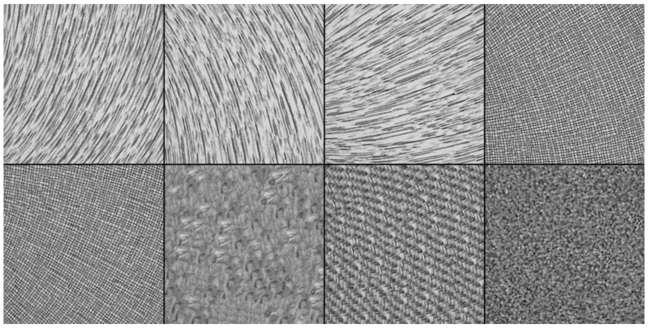

Figure 9: Convex ellipsoids in orthographic projection viewed obliquely.

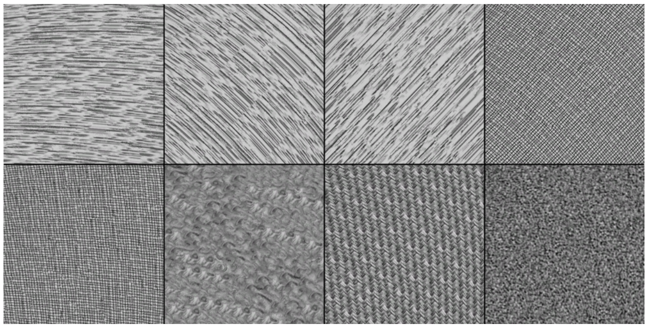

Figure 10: Convex cylinders in orthographic projection viewed obliquely. (The concave set of images was nearly identical in appearance, except for differences that could be eliminated by image rotation.)

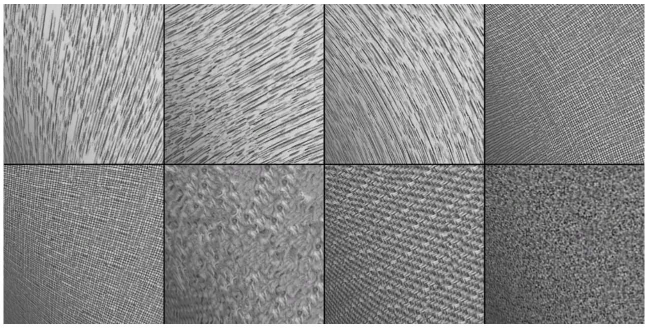

Figure 10: Saddle surfaces in orthographic projection viewed obliquely.

observers.) Five of the observers completed the full set of trials, which included images from all conditions presented in random order. The three other observers completed a reduced version of the experiment, which involved only the perspective projection condition. All observers had normal or corrected-to-normal visual acuity and had no known visual abnormalities. They were given a short training session in order to gain familiarity with the categorization. To avoid introducing biases, we relied on written instructions, which the subjects had to read before they began. Because of the simple nature of the task, that strategy worked well for this situation.

3.1.3 Training

Prior to the experiment, participants were asked to visually and haptically inspect a set of actual 3D hand-sculpted clay objects representing all possible combinations of shape category and orientation. The shape information was labeled on the surfaces in pencil, to avoid the necessity of any verbal explanation that might inadvertently bias the observers to give special attention to the principal directions. Participants were allowed as much time as they felt needed to become familiar with the objects. Figure 12 shows a snapshot of the training surfaces.

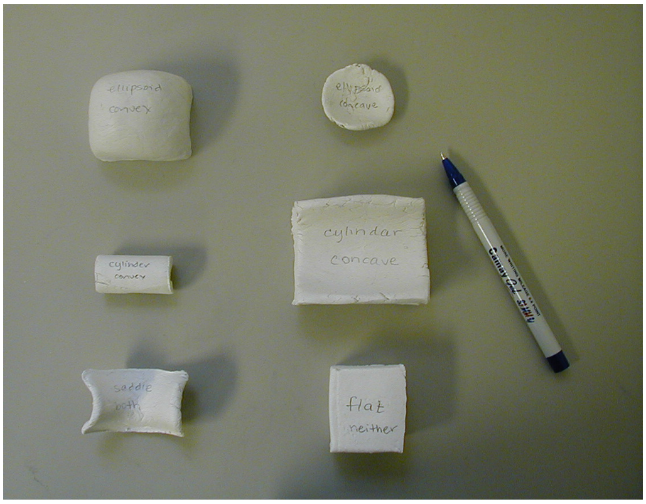

Figure 12: The training surfaces.

3.1.4 Task

The participants were asked to perform two 4AFC (Four-Alternative-Forced-Choice) tasks per trial. One task was to categorize the shape of the surface shown in the trial image as ellipsoid, cylinder, saddle, or flat. The other task was to identify the surface orientation as convex, concave, both, or neither. To conserve participants' effort, we set up the interface to automatically select the shape type 'neither' if the observer chose the shape class 'flat'. The observers recorded their choices by pressing a button on the screen using the mouse. The buttons for the shape category were located vertically above the buttons for the surface orientation. Only after participants had selected options in each grouping could they move on to the next trial. No feedback was given during the experiment. Subjects were shown a white noise image between trials.

3.2 Findings

The charts in figures 13-18 summarize the findings of our experiment. The error bars reflect the boundaries of the 95% confidence interval. Our main findings were: 1) confirmation of the hypothesis that texture type has a significant effect on shape perception, with indications that shape categorization accuracy is generally highest under the principal direction grid texture condition; 2) confirmation of the hypothesis that shape perception is facilitated under conditions of perspective, as opposed to orthographic, projection, overall, but with indications that shape classification rates remain decent, for many textures, in orthographic images when the view direction is oblique (a more generic condition than straight-on viewing); 3) confirmation that convex *orientations* cannot be distinguished from concave orientations under conditions of orthographic projection, regardless of viewing direction or texture type.

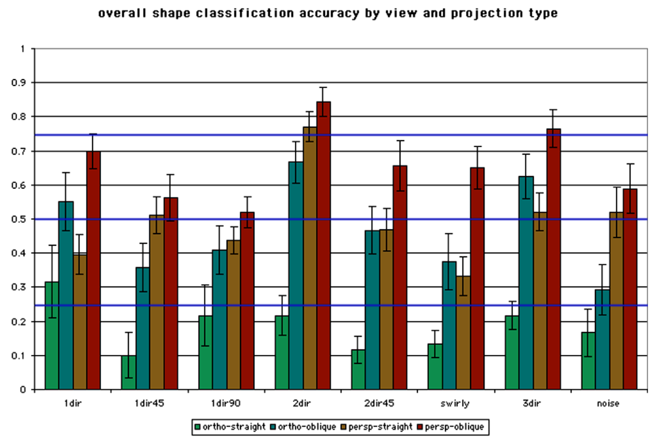

Figure 13: Summary chart showing overall rates of shape classification accuracy, averaged over the three surface shape conditions: ellipsoid, cylinder and saddle, broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique). Results are best in the 2dir texture condition, where accuracy is reliably above 70% under perspective projection for both oblique and straight views.

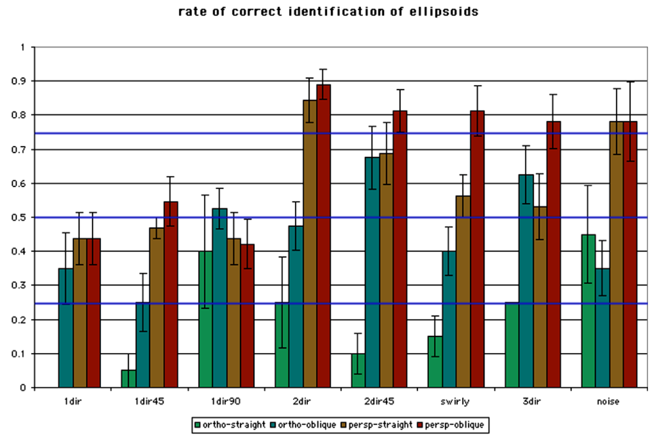

Figure 14: Rates of correct shape classification of ellipsoids (with results pooled over the convex and concave orientations) broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique). Again, performance is best in the 2dir texture condition, with accuracy reliably above 80% under perspective oblique viewing, and reliably above 75% under perspective straight viewing.

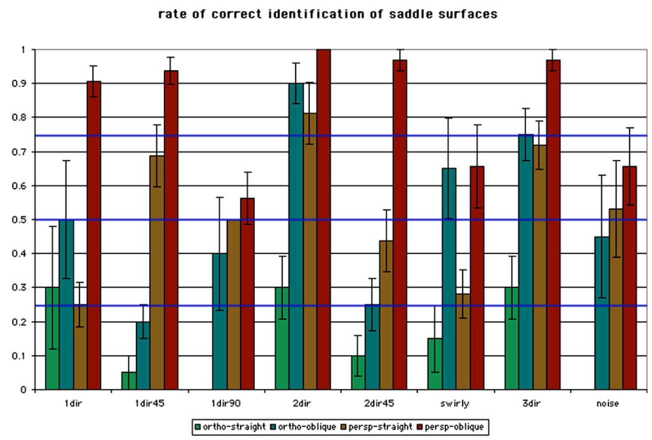

Figure 15: Rates of correct shape classification of saddle surfaces, broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique). Accuracy is a perfect 100% for the 2dir texture in the perspective/oblique viewing condition, and reliably above 75% under perspective/oblique viewing for all textures except swirly, noise, and 1dir90. Accuracy falls dramatically under the other viewing conditions except in the case of the 2dir pattern, where it remains reliably above 70% across both the perspective/straight and orthographic/oblique conditions.

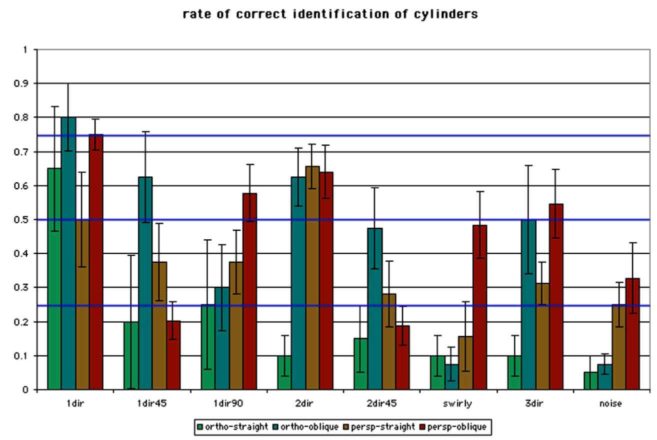

Figure 16: Rates of correct shape classification of cylinders (with results pooled over the convex and concave orientations) broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique). Accuracy rates are lower, overall, than in the cases of the doubly curved surface patches. Results are reliably above 50% only for the first principal direction and 2dir textures. However, further analysis turned up a particularly high rate of false positives for the cylinder classification in the case of the 1dir textures (observers guessing 'cylinder' when the surfaces were actually flat).

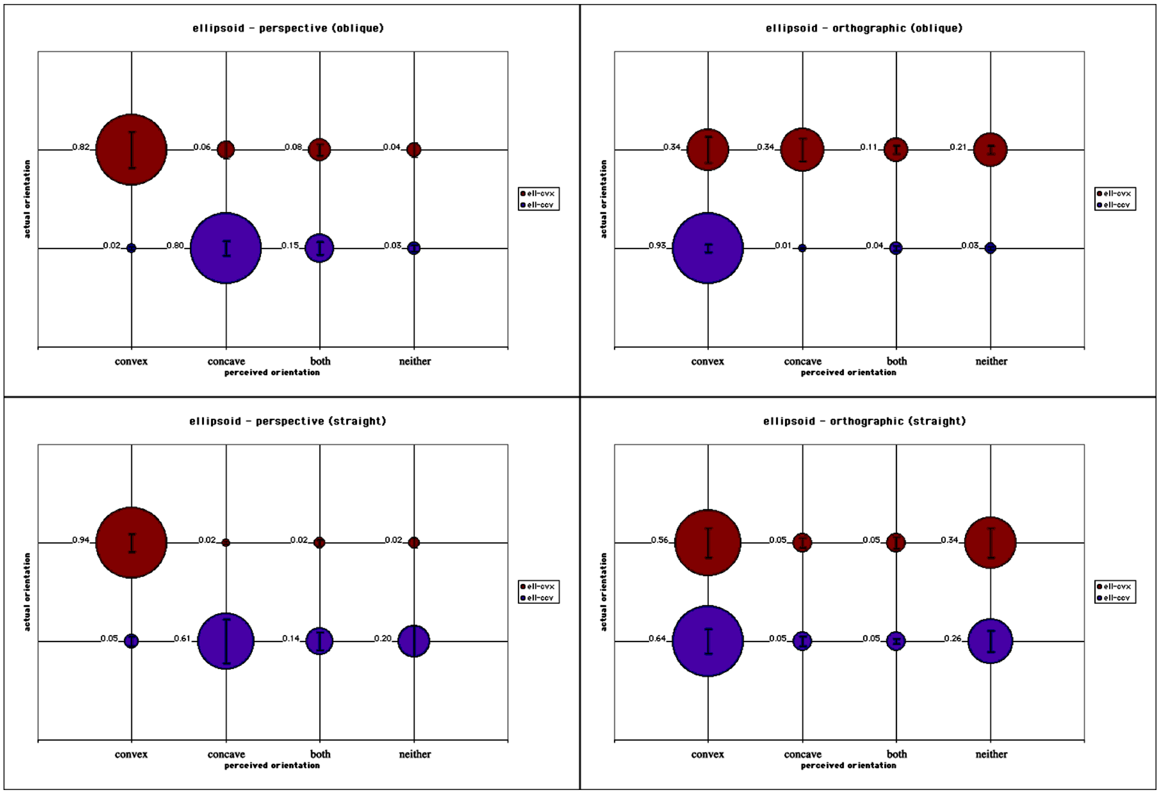

Figure 17: Rates of correct classification of ellipsoid orientation, using results pooled over all texture types. The charts are arranged horizontally by projection type (perspective then orthographic), and vertically by viewpoint (oblique then straight-on). The columns within each chart specify the response frequency in each orientation category (convex, concave, both, neither) for surfaces whose actual orientation is either convex (top row) or concave (bottom row). Results are excellent in the perspective cases but abysmal under orthographic projection, where not only is there confusion between convex and concave but also a greater tendency to perceive the shapes as flat. The asymmetry in the chart on the upper right between the convex and concave conditions is most likely an artifact due to the inconsistency of the oblique view directions used in these two cases.

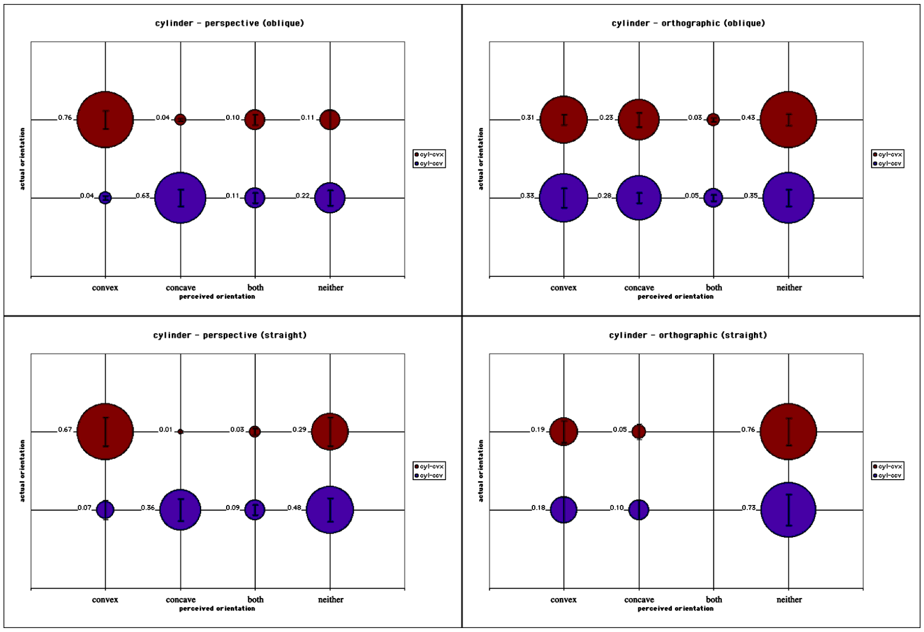

Figure 18: Rates of correct classification of cylinder orientation (convex/concave/both/neither), using results pooled over all texture types. Results basically follow the same pattern as with the ellipsoids.

4. DISCUSSION

Looking more closely at the details of all of the results, we observed several things. 1) The noise and swirly textures permitted reliable correct discrimination only in the case of the ellipsoid shape. 2) The diagonal textures caused particular problems for the cylinder recognition, but not for the recognition of saddle and flat surfaces. Results for ellipsoids was mixed: the 45-degree rotated 2dir texture worked well, but 1dir45 did not. 3) The textures that followed just one of the principal directions caused ellipsoids and saddles to be perceived as cylinders often enough that discrimination was unreliable. They did not appear to interfere with the recognition of cylinders as cylinders, but at the same time there were also a high number of cylinder responses when the surfaces were actually flat (possibly indicative of a response bias towards indicating a perception of curvature even when no curvature was in fact present.) 4) The three-directional texture worked well for the flat and saddle surfaces, but not as well for the ellipsoid or cylinder. It is possible that the presence of the diagonal component in the 3dir texture interfered with shape perception in the cylinder case in the same way that it did in the 1dir45 case, which didn't include the pdir components. The reason for the relatively poorer performance with the 3dir texture in the ellipsoid case is not immediately clear. One potentially complicating factor is that this texture, being drawn from a photograph, was slightly less rigidly regular as the other directional patterns, which were created artificially. It could be that people were using a slightly different strategy in this case than in the others.

The conservative conclusion is that adding information along more directions in addition to the principal direction is not 100% safe - it does not seem to help, and it may in fact hurt.

The main results from this experiment are:

1) Shape classification accuracy rates were highest, overall, under the condition of the principal direction grid texturing (2dir). Using an ANOVA analysis and looking at pairwise comparisons using Tukey's HSD test we found that there was no case in which shape classification was significantly worse with the 2dir texture than with any other texture. In many, but not all, cases, classification accuracy was significantly better with the 2dir texture than with the other patterns. The 2dir texture was the only texture that gave consistently reliable performance, under conditions where accuracy was possible.

2) High rates of shape classification accuracy were achieved under some texture conditions despite the use of orthographic projection, as long as the surface was viewed from an oblique vantage point. In particular, shape classification accuracy was as good with the oblique/orthographic viewing as with the perspective/straight viewing, for many texture types. Shape classification accuracy was abysmal under orthographic projection when the view direction was directly head-on to the surface, due to the loss of critical surface orientation cues in the stimuli because of the non-generic viewing condition.

Other observations we can make are:

1) In an ANOVA analysis considering results across all texture types, projection conditions and viewing configurations, we found that shape classification accuracy was significantly higher under perspective projection than under orthographic projection, and in an oblique view than in a head-on view.

2) With the orthographic projection, elliptical patches were equally likely (in the oblique view) or more likely (in the straight view) to be perceived as cylindrical, while hyperbolic patches were equally likely to be perceived as cylindrical, and, in the straight view, more likely to be misperceived as flat.

3) Cylindrical patches were reliably more likely to be recognized as cylindrical than as some other shape only under conditions of perspective projection and an oblique view direction. In all other conditions they were equally likely (perspective/straight and

orthographic/oblique) or more likely (orthographic/straight) to be misidentified as flat.

Considering recognition rates for each texture and each shape separately, we find that correct recognition rates for elliptic patches are reliably above 80% only for the principal direction grid texture, and even then, only under perspective projection and an oblique view. Correct recognition rates are reliably above 60% for elliptical patches with the principal direction grid texture, the 45-degree rotated principal direction grid texture, and the noise texture, under perspective projection for both straight and oblique viewing conditions, and for the swirly texture and the three-directional texture under perspective projection and an oblique view. They are not reliably above 50% for any of the one-directional textures, regardless of projection type or view.

Recognition rates for hyperbolic patches are reliably above 75%, under perspective projection and with an oblique view, for the principal direction grid texture, the 45-degree rotated principal direction grid texture, the three-directional texture (following both principal directions and a also a third direction, diagonal to the principal directions), the first principal direction texture, and the 45 degree rotated first principal direction texture. Recognition rates are also reliably above 75% for the principal direction grid texture under conditions of orthographic projection but with an oblique view. Correct recognition rates are reliably above 50% for hyperbolic patches for the principal direction grid texture, the three-directional texture, and the 45 degree rotated first principal direction texture under perspective projection for a straight-on view, and for the three-directional texture under conditions of orthographic projection but with an oblique view. They are not reliably above 50% for any texture under orthographic projection and head-on viewing, nor for the second principal direction texture under any condition of projection and viewing.

In an ANOVA analysis of the perspective data alone, we found that the oblique viewpoint was significantly better ($p<0.05$) than the straight viewpoint, for both shape classification and surface orientation judgments, overall (pooled over texture type and surface type). This was true individually, also, for saddles (shape and orientation), ellipsoids (shape only), and cylinders (shape and orientation).

Finally, we noticed that the textures containing diagonal elements were especially effective at reliably indicating flat surfaces, in perspective projection, pooling over straight and oblique views.

5. CONCLUSIONS

These are some of the first studies that have systematically investigated the effects on shape perception, measured by shape classification accuracy, of using a principal direction vs. non-principal direction texture pattern, in the case of doubly curved surfaces. In this study, we have found that shape classification is never significantly worse with the principal direction grid texture (2dir) than with any other texture pattern, according to pairwise comparisons. Likewise, we have found that rates of correct surface orientation judgments (convex/concave) are never significantly lower under the 2dir texture condition than with any other pattern. These results are important because they indicate that in general, practical situations, when you can’t be sure what shape your surfaces will have at any particular point, you won’t go wrong by choosing a principal direction texture.

6. ACKNOWLEGMENTS

This work was supported by an NSF Presidential Early Career Award for Scientists and Engineers (ACI-9875368). We are grateful to the Statistical Consulting Clinic, which is supported by funding from the Minnesota Agricultural Experiment Station, for assistance with the experimental analysis. The texture synthesis and image rendering software was written by Gabriele Gorla, with support from a University of Minnesota Grant-in-Aid of Research, Scholarship and Artistry, who was co-advised by Guillermo Sapiro. We are also indebted to our observers AJ, TU, JM, II and TF for their dedicated and conscientious efforts.

References

[1] Erens, R.G.F., Kappers, A., Koenderink, J. (1993). Perception of local shape from shading, Perception & Psychophysics, **54**(2): 145-156.

[2] Gorla, G., Interrante, V. and Sapiro, G. (2003). Texture Synthesis for 3D Shape Representation, IEEE Transactions on Visualization and Computer Graphics, **9**(4), Oct-Dec 2003.

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

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

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

[6] Kim. S., Hagh-Shenas, S. and Interrante, V. (2003). Showing Shape with Texture: Two Directions Seem Better than One, Human Vision and Electronic Imaging VIII, SPIE **5007**, pp. 332-339.

[7] Knill, D. (2001). Contour into Texture: Information Content of Surface Contours and Texture Flow, Journal of the Optical Society of America, A, (1), January 2001, pp. 12-35.

[8] Koenderink, J.J., Doorn, A. and Kappers, A. (1992). Surface Perception in Pictures, *Perception*, **52**, pp. 487–496.

[9] Li, A. and Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. Vision Research, **40**(2): 217-242.

[10] Li, A. and Zaidi, Q. (2001). Information limitations in the perception of shape from texture. Vision Research, **41**(22): 2927-2942.

[11] Mamassian, P. and Kersten, D. (1996). Illumination, shading and perception of local orientation. Vision Research, **36**(15): 2351-2367.

[12] Mamassian, P. and Landy, M. (1998). Observer Biases in the 3D Interpretation of Line Drawings, Vision Research, **38**(18): 2817–2832.

[13] Rosenholtz, R., and Malik, J. (1997). Surface orientation from texture: Isotropy or homogeneity (or both)? Vision Research, **37**(16): 2283-2293.

[14] Stevens, K. (1981). The Information Content of Texture Gradients, *Biological Cybernetics*, **42**, pp. 95–105.

[15] Todd, J.T. and Oomes, A.H.J. (2002). Generic and non-generic conditions for the perception of the surface shape from texture. Vision Research, **42**(7): 837–850.

[16] Todd, J., Norman, F., Koenderink, J. and Kappers, A. (1997). Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. Perception, **26**, pp. 807-822.

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