Conveying Shape with Texture: experimental investigations of texture's effects on shape categorization judgments

Sunghee Kim, Haleh Hagh-Shenas, and Victoria Interrante, *Senior Member, IEEE*

Abstract—In this paper we describe the results of two comprehensive controlled observer experiments intended to yield insight into the following question: 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? We begin by reviewing the results of our initial study in this series, which were presented at the 2003 IEEE Symposium on Information Visualization, and offer an expanded analysis of those findings. We continue by presenting the results of a follow-on study, in which we sought to more specifically investigate the separate and combined influences on shape perception of particular texture components, with the goal of obtaining a clearer view of their potential information carrying capacities. In each study we investigated observers' ability to identify the intrinsic shape category of a surface patch (elliptical, hyperbolic, cylindrical or flat) and its extrinsic surface orientation (convex, concave, both or neither). In our first study we compared performance under 8 different texture type conditions, plus 2 projection conditions (perspective or orthographic) and 2 viewing conditions (head-on or oblique). In this study we found that: 1) shape perception was better facilitated, in general, by the bi-directional 'principal direction grid' pattern than by any of the 7 other patterns tested; 2) shape type classification accuracy remained high under the orthographic projection condition for some texture types when the viewpoint was oblique; 3) perspective projection was required for accurate surface orientation classification; and 4) shape classification accuracy was higher when the surface patches were oriented at a (generic) oblique angle to the line of sight than when they were oriented (in a non-generic pose) to face the viewpoint shtraight-on. In our second study, we we compared performance under 8 new texture type conditions, redesigned to facilitate gathering insight into the cumulative effects of specific individual directional components in a wider variety of multi-directional texture patterns. In this follow-on study we found that shape classification accuracy was equivalently good under a variety of test patterns that included components following either the first or first and second principal directions, in addition to other directions, suggesting that a principal direction grid texture is not the only possible 'best option' for enhancing shape representation.

Index Terms— Three-Dimensional Graphics and Realism, Vision and Scene Understanding.
Additional Keywords— Shape Perception, Shape Representation, Texture, Principal Directions.

1 Introduction

OUR goal, as visualization designers, 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 and definition 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 structures defined by complicated, smoothly curving surfaces in their data. The most common practice in rendering objects for visualization purposes 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, as hinted in figure 1, smooth shading is not optimal for all purposes and in particular is not optimal for shape representation. Research in shape perception has consistently shown that shape understanding can be facilitated by the presence of the right kinds of surface texture. Unfortunately, existing theories do not yet provide sufficient guidance to tell us how exactly to specify a texture pattern that can most effectively and accurately convey surface shape.

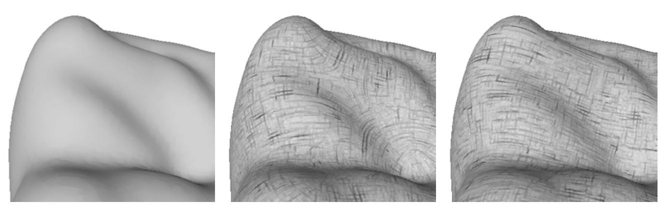

Figure 1: A close-up view of the top portion of a tooth dataset, depicted, from left to right, with: no texture, with an orthogonal texture pattern following smoothed principal directions, and with an orthogonal texture pattern following constant uniform directions in object space.

Over the past several years, we have conducted a series of experiments [1, 2, 3] investigating the impacts of various characteristics of surface texture patterns on shape perception. In these studies, we have found that observers' judg-

- *S.K. is a PhD student in the Department of Computer Science, University of Minnesota. e-mail: skim@cs.umn.edu.*
- *H.H. is a PhD student in the Department of Computer Science, University of Minnesota. e-mail: haleh@cs.umn.edu.*
- *V.I. is an Aassistant Professor in the Department of Computer Science, University of Minnesota. e-mail: interran@cs.umn.edu..*

ments of local surface orientation, made with a surface attitude probe, are most accurate under conditions of anisotropic texturing when the direction(s) of the texture anisotropy are 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 varies over the surface in a manner unrelated to the surface geometry. However, many questions remain.

In this paper we present the findings of our most recent experiments in the evaluation of alternative texture patterns for shape representation. But before describing the current work we provide a brief review of background material and give an overview of previous work in shape perception from shading and texture.

2 Background

Researchers have long known that our ability to accurately perceive a surface's shape can be facilitated — or impeded — by characteristics of shading and texture, and many studies have been undertaken over the years to elucidate both the nature of these effects and their possible neural or developmental origins.

Observation tells us that shading clearly plays an important role in conveying information about the shape structure of a surface, and research in shape from shading suggests that shape perception is best facilitated under conditions of "natural", overhead, oblique illumination. 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. Erens *et. al* [4] found that observers were unable to reliably disambiguate elliptic (egg-shaped) from hyperbolic (saddle-shaped) surface patches from local views of smoothly shaded quadric surface patches in which contour cues were unavailable. Mamassian and Kersten [5] found that when contour cues are available in images of smoothly curving untextured surfaces, observers appear to rely upon them exclusively when asked to estimate local surface orientation at points in adjacent interior patches, basically ignoring any shape cues provided by the shading.

Numerous studies have found that shape perception can be enhanced by the use of an appropriate surface texture pattern in addition to shading (e.g. Todd *et. al* [6]). Cumming *et. al* [7] showed that providing appropriate texture cues significantly enhances accurate shape perception, under conditions of stereo viewing and in the presence of visible surface contours, a result also confirmed in later studies by our group [1]. Since it has been well established that shape perception can be enhanced by the addition of appropriate texture under generic viewing and shading conditions, it is without loss of generality that we restrict our investigations in the present study to the particular case of a tightly controlled set of stimuli in which shading and contour cues do not play a significant role in indicating shape information.

3 Previous Work

Because reliable computer graphic techniques for applying arbitrary given texture pattern to an arbitrary doubly curved surface have only recently been developed, e.g. [8], research on shape perception from texture has generally been somewhat restricted, either to developable surfaces [e.g. 9,10] (surfaces which can be rolled out to lie flat on a plane), to patterns projected onto surfaces from a particular direction [e.g. 11], or to solid textures [e.g. 6] (whose features are generally independent of the shape of surface that is carved out of them). Although there are many open questions that still exist about the impact of surface texture on shape perception, many important insights have also already been achieved.

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. [12] but also upon how the pattern is laid down over the surface [1,10,13].

Stevens [14] and Mamassian and Landy [15] have suggested that observers may be biased toward interpreting lines on surfaces as if they were following the principal directions. Li and Zaidi [10,13] have found that two conditions are necessary for the perception of 3D shape from texture: 1) when the surface is viewed straight-on, 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 measure shape perception, discriminating which of two adjacent points is more distant, actually only provides coarse information about the perceived direction of surface slant. This is sufficient for determining whether observers can differentiate convexities from concavities, but does not capture all of the information that we might like to know about shape perception.

Other researchers have downplayed 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, Todd and colleagues [11,16] 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 pattern itself lacks significant energy in any particular direction. Todd *et. al* furthermore 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: 1) 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]; 2) 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 [e.g. 10]; 3) identification of the quadrant in which two surfaces differ in shape [e.g. 2]; and 4) identification of the shape category of a surface patch [4].

In our own previous work we have found indications that surface attitude judgments are significantly more accurate [1] and surface shape discrimination thresholds significantly reduced [2] 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 the current studies we sought to compare alternative principal direction oriented patterns. First, 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 pattern aligned only with the first or the second principal direction. In a previous study [3] we had found indications that this might be the case, but our results were below significance at the 95% level. In the current experiments we also sought to investigate the impact on shape perception of employing patterns containing elements systematically oriented at an oblique angle to the principal direction(s). Such textures might implicitly encode the principal directions, but the eye would be drawn by these texture to follow lines over the surface that are significantly different from the principal directions, and the possibility would exist that observers might interpret the information provided by these oblique lines as if it were provided by lines oriented in the principal directions even though it was not. Finally, we were also interested in testing the claim by Li and Zaidi that accurate shape perception requires perspective projection. Like Todd, we had accumulated some anecdotal evidence that shape perception might be still possible even under conditions of orthographic projection, and we were interested to pursue this question further.

4 EXPERIMENT 1

In our first experiment, we set out to answer the following 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.

4.1 Method

4.1.1 Design

In this study we used a within-subjects design, in which participants were asked to perform two four-alternative-forced-choice tasks. The first task was to identify the intrinsic shape of a surface patch shown in an image as one of the following four types: ellipsoid, cylinder, saddle, or flat. The second task was to identify the orientation of the surface as being convex or concave (or 'both', if the shape type was saddle, or 'neither' if the shape type was flat). We used eight different texture patterns, described in the next subsection, two different viewing conditions (straight-on and oblique), and two different projection conditions: perspective and orthographic. To control for orientation dependent effects we also rotated each image stimulus in the plane over repeated trials, using two rotations – 0° and 90° – for the axisymmetric straight-view images, and four rotations – 0°, 90°, 180° and 270° – for the oblique view images.

4.1.2 Stimuli

Our surface stimuli for this study consisted of simple quadric patches. We constructed the patches as height fields defined by the parametric equations for each of the following shape categories: ellipsoid, elliptical cylinder, saddle, and flat, being careful to use consistent coefficients between shape types to ensure that the curvatures of each patch were as consistent as possible. We determined the first and second principal directions analytically at each of the vertices in the surface meshes using the parametric formulas.

We used the algorithm developed by Gorla et al. [8] to synthesize each of the eight sample patterns shown in figure 2 over each of our test surfaces in such a way that the final surface texture was everywhere seamless and nondistorted, and locally aligned with the principal direction coordinate frame at a per-pixel level. Because we were interested in looking at both the convex and concave orientations of each patch, we textured both sides of each surface. The eight patterns that we chose can be described as follows: 1-directional, (elongated in the first principal direction), 1-directional rotated clock-wise 45 degrees (diagonal to the first principal direction), 1-directional rotated clockwise 90 degrees (to be aligned with the second principal direction), 2-directional (indicating both the first and second principal directions), 2-directional rotated clock-wise 45 degrees, 3-directional (encoding the first and second principal directions plus a diagonal direction), swirly (turning in the surface) and filtered white noise (isotropic).

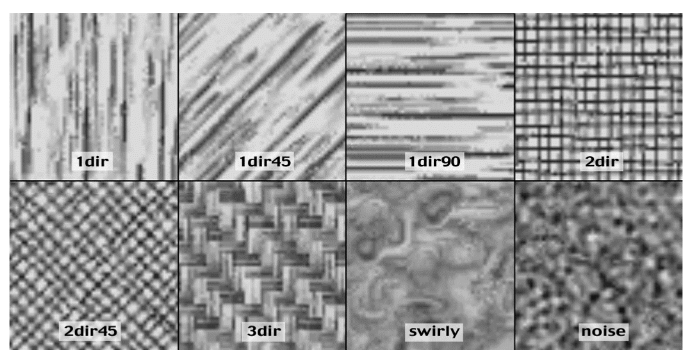

Figure 2: The sample texture patterns used in the study. When applied to the surfaces, the vertical direction in each pattern is aligned with the first principal direction over the surface.

The surface stimuli were imaged using two different fixed viewing directions: head-on and oblique, and two different projection types: perspective and orthographic.

For each viewing configuration, we determined a fixed camera position from which we could obtain a snapshot of each surface patch in which no edges were visible, yet in which a sufficiently large portion of the surface interior remained visible. We took great pains to define viewing parameters that were as nearly identical as possible across the different surface types, and between the convex and concave views of each surface. Unfortunately it was not possible to find a single camera position that worked for all situations, and an exception was necessary in the case of the oblique views of the convex and concave ellipsoids. Figure 3 shows thumbnail images of some of the surface stimuli.

All surfaces were subtly shaded using a standard Phong illumination model with an oblique directional light source.

4.1.3 Procedure

Participants recorded their choices by pressing a button on the screen interface using the mouse. The buttons for the shape type were located directly above the buttons for the surface orientation. We set up the interface to automatically select the orientation option 'both' if the observer chose the shape type 'saddle', and the orientation option 'neither' if they chose 'flat'. The interface required participants to

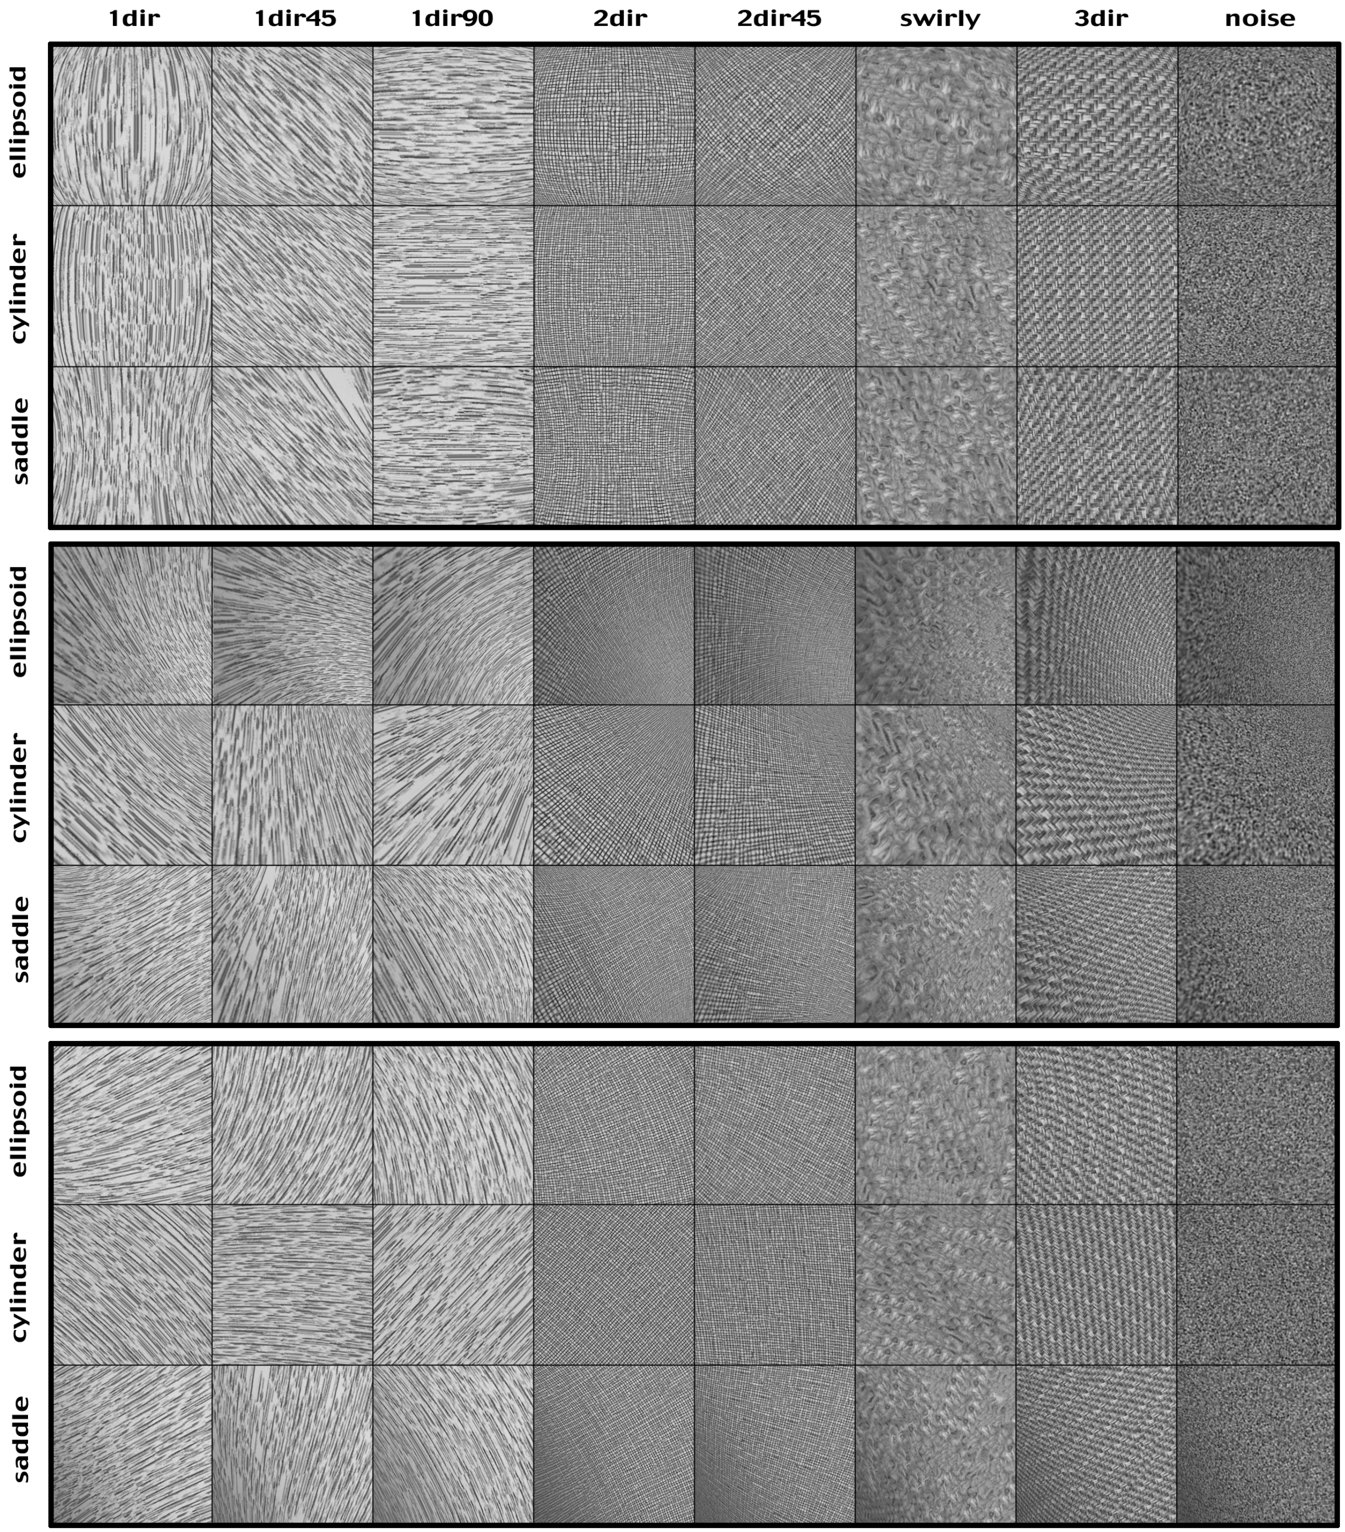

Figure 3: Thumbnail images of some of the 592 sample stimuli used in the experiment. Top: convex orientation, straight-on view and perspective projection; Middle: concave orientation, oblique view and perspective projection; Bottom: convex orientation, oblique view, orthographic projection.

respond to each question before moving on to the next trial, and did not allow returning to previous trials. No feedback was given. There were 592 trials in total, and participants were shown a white noise image between trials.

The images were displayed 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, and rest breaks were enforced at regular intervals.

A total of 8 observers – 5 males and 3 females, ranging in age from 17-50 – participated in the first 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. Three were members of our research team. (We noticed no significant differences in patterns of performance between these different categories of 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. To avoid inadvertently introducing biases, we relied on written instructions, which we had our participants read, along with their consent form, before they began the study. Because of the simple nature of the task, that strategy worked well for this situation.

4.1.4 Training

Prior to the experiment, participants were asked to visually and haptically inspect a set of 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 have inadvertently biased the observers to give special attention to the principal directions. Participants were allowed as much time as they needed to become familiar with the definitions of the four shape categories and surface orientations, and they were free to refer back to the models at any time during the experiment. Figure 4 shows a snapshot of the training surfaces.

Figure 4: The training surfaces.

4.2 Findings

The charts in figures 5-8 summarize some of the main findings of our first experiment. The error bars in these charts indicate the standard errors of the means.

We performed a 5-way analysis of variance on the results, looking at the true positive rate for correct shape identification as a function of subject x texture_type x shape_type x projection x view. We found significant main effects for all five variables: subject ($F_{4,4640}$=9.969, p<<0.001), texture_type ($F_{7,28}$=16.62, p<<0.001), shape_type ($F_{2,8}$=5.187, p<0.05), projection ($F_{1,4}$=53.27, p<0.01) and view ($F_{1,4}$=79.27, p<0.01). What this means is that the rate of correct shape identification was significantly different for different observers, different textures, different surface types, different projections and different views.

We also found significant two-way interactions between subject and {texture_type, shape_type, projection and view}, as well as between texture_type and shape_type ($F_{14,56}$=5.781, p<<0.001), texture_type and projection ($F_{7,28}$=3.135, p<0.05), and shape_type and projection ($F_{2,8}$=6.234, p<0.05), as well as many 3-way and 4-way interactions between various combinations of variables. What this means is that certain people systematically performed better than others under specific conditions, and that rates of correct shape identification were systematically different under particular combinations of variables.

To gain more insight into the nature of the various significant main effects, we performed a series of Tukey HSD ("honest significant difference") post-hoc analyses. The results of the main analysis are summarized in figure 9, which shows the coefficients of the effects of the separate texture types, color coded to indicate areas of overlap in significance. We can see that when results are pooled across all three meaningful shape types, shape categorization accuracy is significantly better with the 2dir texture pattern than with any of the other patterns tested. Not shown in chart form are the results of the HSD analyses on projection and view, which showed that shape classification accuracy is significantly (31%) better under perspective than orthographic projection (coeffs = {-0.157 orthographic, 0.157 perspective}, p<<0.001), and significantly (22%) better for oblique than straight views (coefficients= {-0.11 straight, 0.11 oblique}, p<<0.001).

Looking separately at the results for the ellipsoid, saddle and cylinder stimuli, we can gain some deeper insight into the details of the effects of texture on shape categorization accuracy in these cases. In the post-hoc analysis we found that shape classification accuracy was significantly better (p<0.05) for the doubly curved (generic) surfaces (ellipsoids and saddles) than for the singly curved (non-generic) surfaces (cylinders). We ran a 3-way ANOVA separately on the subsets of ellipsoid and cylinder stimuli, looking at shape classification accuracy as a function of subject x texture x orientation (convex vs. concave), and found a significant main effect of orientation in the case of the ellipsoids ($F_{1,7}$=7.936, p<0.05) but not cylinders ($F_{1,7}$=2.25, p=0.1773). The post-hoc analysis showed that in the cases of the ellipsoids the accuracy of shape category judgments was significantly better with the convex surface orientations than with the concave orientations.

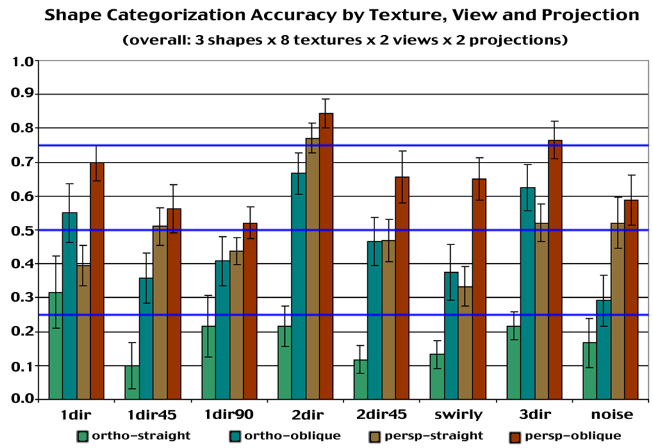

Figure 5: Summary chart showing overall rates of shape classification accuracy, averaged over the three surface shape conditions: ellipsoid, cylinder and saddle, broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique).

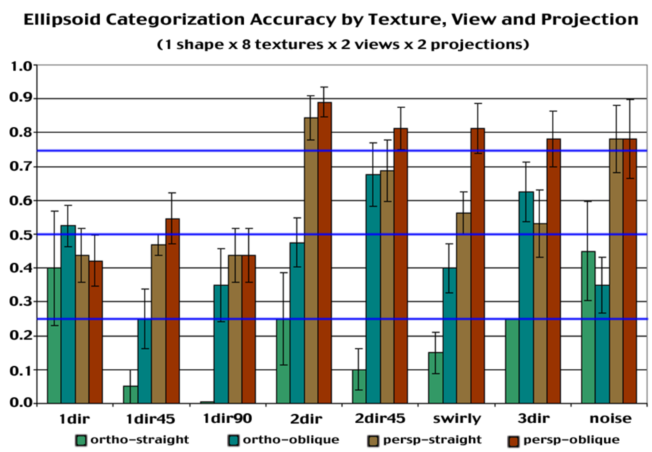

Figure 6: Rates of correct shape classification of ellipsoids (with results pooled over the convex and concave orientations), broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique).

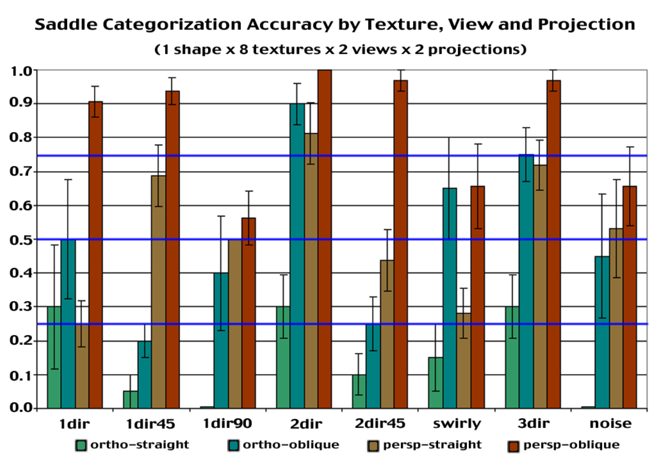

Figure 7: Rates of correct shape classification of saddle surfaces, broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique).

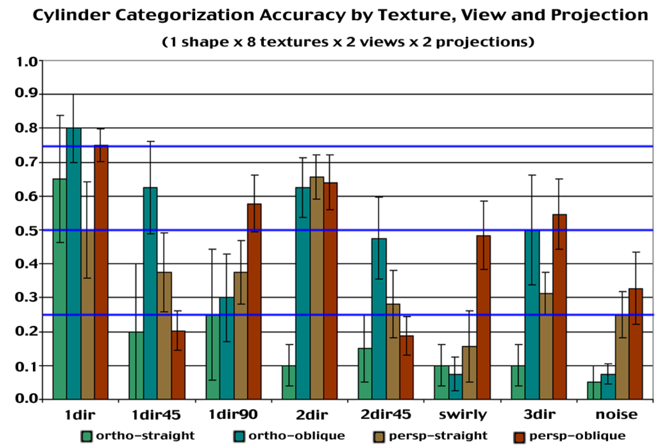

Figure 8: Rates of correct shape classification of cylinders (with results pooled over the convex and concave orientations) broken down by texture type, projection type (perspective or orthographic), and viewpoint (straight-on or oblique).

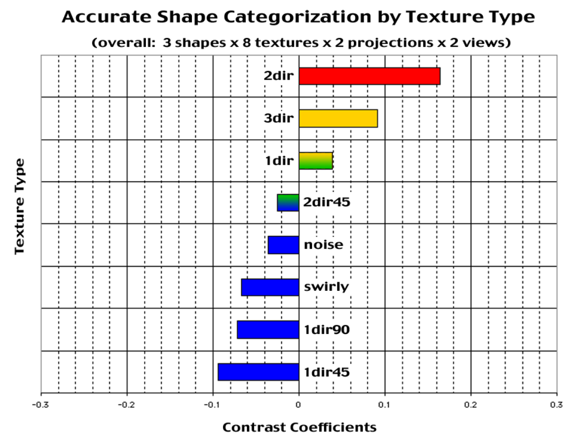

Figure 9: Contrast coefficients from the Tukey HSD comparison of the effects of texture type on shape classification accuracy. Components containing the same color are statistically similar at $p<0.05$.

It can be seen in figures 5-9 that accurate shape classification rates remained at levels significantly above chance under some of the more impoverished, non-generic and non-ecological projection and viewing conditions, particularly in the cases when the 2dir pattern was used. However, finding that shape categorization was significantly better, overall, for the stimuli imaged under the generic conditions of perspective projection and oblique viewing, we decided that for greatest generality we would focus our further efforts on this subset of the data in particular.

Figure 10 shows the results of the application of the Tukey HSD analysis on the effects of texture type to the data from the ellipsoid, cylinder and saddle shape classes separately. It is clear that different texture components are critical in facilitating shape understanding in each of these three cases. In particular: accurate identification of ellipsoid stimuli is particulary poor under the one-directional texture pattern conditions; accurate identification of saddle surfaces is better under any of the tested directional – as opposed to non-directional – pattern conditions, with the exception of 1dir90; and accurate perception of cylinders is particularly impeded when a diagonal pattern is applied.

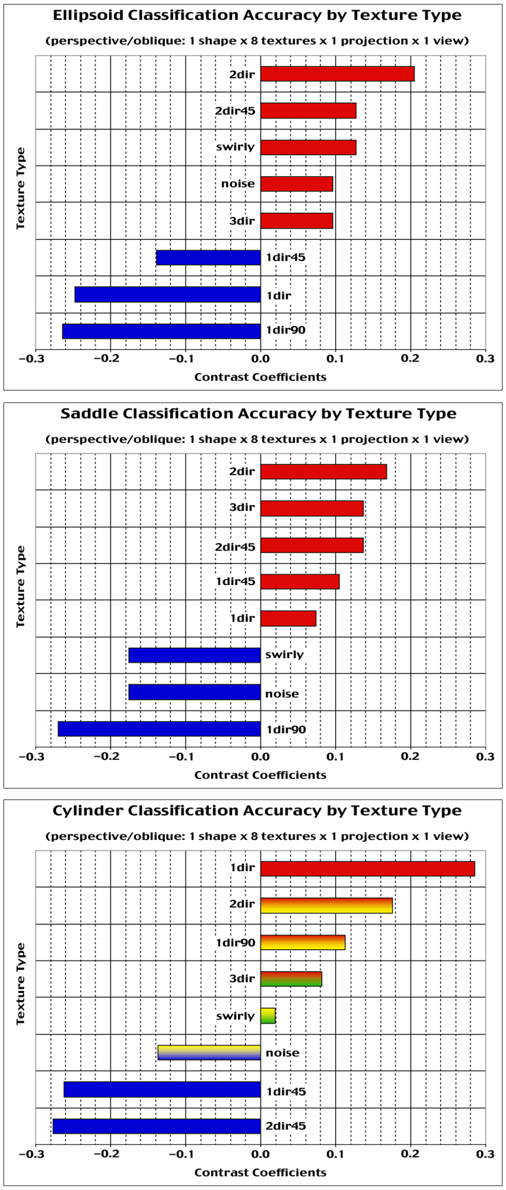

Figure 10: Contrast coefficients from the Tukey HSD comparison of the effects of texture type on shape classification accuracy. Limited to the perspective/oblique stimuli and broken down by shape category. Color coding indicates statistical similarity (p<0.05).

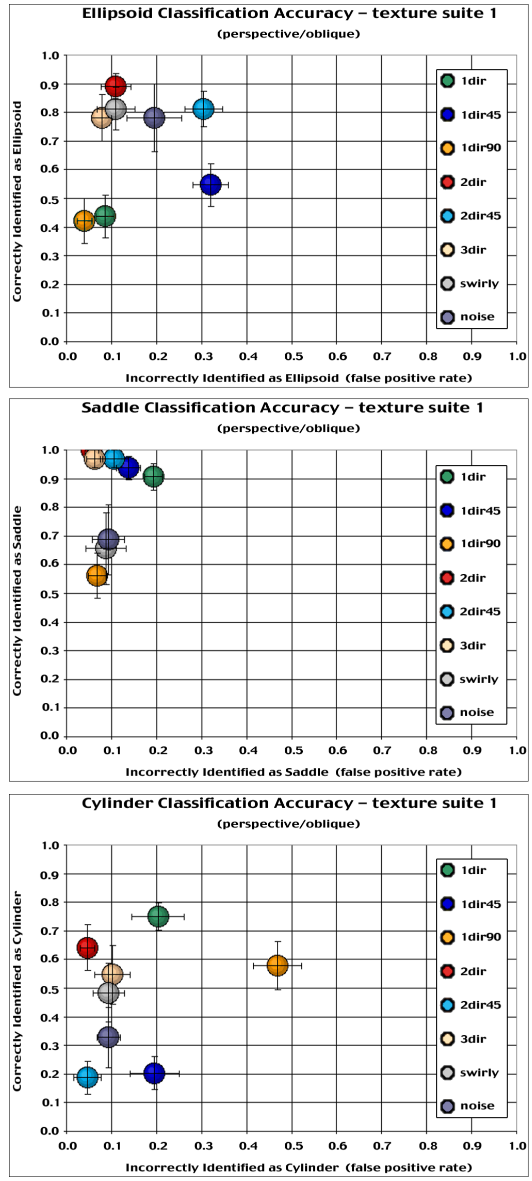

Figure 11: An alternative representation of shape classification accuracy, under conditions of perspective projection and oblique viewing, broken down by texture type (with results pooled over the convex and concave orientations).

Further insight into the experimental results can be found through an examination of the rates of false positive as well as true positive responses. Figure 11 plots the true and false positive rates for the perspective/oblique stimuli, broken down separately by shape category. Ideal classification under a particular texture type condition would result in a point at (0,1), in the upper left corner of these graphs, corre sponding to a situation in which there is a 100% true positive classification rate associated with that surface category and texture type and a 0% false positive classification rate. If observers are "over guessing" a particular shape category when a particular texture pattern is used — a situation that might arise, for example, with texture patterns that mask shape information and hence invite abundant erroneous indications of the shape type 'flat' — that event would show up in these charts as an elevated rate of false positive responses, which would place the point encoding the performance under that texture condition farther horizontally to the right in the chart. As in the previous set of charts, the vertical height of each data point the image is determined by the magnitude of the true positive response rate.

In figure 11, we can clearly see high false positive rates for cylinder classification in the three unidirectional conditions, indicating that observers may be tending to interpret surface curvature components as flat when reliable cues to the curvature in these directions are absent. False positive rates for saddle classification are generally low, except in the case of the 1dir texture, probably due to the misclassification of concave ellipsoids as saddles when the 1dir pattern contains prominent sets of nearly parallel curves of opposing sign that resemble the contour edges of a hyperbola. Finally, false positive rates for ellipsoid classification are relatively highest in the case of patterns that trace out curves diagonal to the principal directions, under which conditions cylindrical patches are most likely to be misperceived as being doubly curved. When a texture pattern is associated with a high false positive rate, it is an indication that use of that pattern can lead to systematically misleading interpretations under certain conditions, which is something that we would like to avoid.

Using the representational approach shown in figure 11, it is easier to see the previously summarized main result, obtained from the post-hoc analysis, that shape classification accuracy, both in the general case and in the particular case of the perspective/oblique view, is often better, and never significantly worse, under the 2dir texture condition than under any of the other texture conditions studied.

To summarize, the main findings thus far are: 1) confirmation of the hypothesis that texture type has a significant effect on shape perception, with indications that shape classification accuracy is generically highest under the principal direction grid texture condition; and 2) confirmation of the hypothesis that shape perception is particularly facilitated under conditions of perspective, as opposed to orthographic, projection, while recognizing that shape classification rates remain well above chance for many textures in orthographic images when the view direction is oblique.

We now discuss the findings of our second set of results from the present experiment. To analyse these results we performed a 6-way ANOVA looking at rates of correct surface orientation identification (convex vs. concave) as a function of subject x texture_type x shape_category x pro jection x view x image_rotation over the ellipsoid and cylinder data. We found significant main effects for shape category ($F_{1,4}$=18.03, p<0.05), projection ($F_{1,4}$=62.58, p<0.01), and image rotation ($F_{3,12}$=5.334, p<0.05), and significant two-way interactions between subject and texture type, texture type and shape category, shape_category and pro jection, and shape_category and view, as well as assorted significant 3-way, 4-way and even 5-way interactions between various variables. High rates of correct surface orientation judgments are associated with use of the strongly directional texture patterns (1dir, 2dir, 2dir45, 3dir and 1dir90 in that order) and orientation judgments are poorest with the non-directional patterns 'swirly' and 'noise'. Surface orientation judgements are significantly more accurate for the ellipsoid shapes than the cylinders (p<0.05; coefficients = {-0.0523 cyl and 0.0523 ell}).

Most significantly, surface orientation judgments critically depend on the use of perspective projection (p<0.05; coefficients = {=0.249 orthographic and 0.249 perspective}). Convex orientations cannot be distinguished from concave orientations under conditions of orthographic projection. This result is clearly illustrated in the charts shown in figure 12. Here we can see that rates of correct classification of surface orientation are excellent when a perspective projection is used but abysmal under orthographic projection, where not only is there confusion between convex and concave orientations but also a greater tendency to perceive the shapes as flat. Please note that the apparent asymmetry between the convex and concave conditions in the chart on the upper right is most likely an artefact of inconsistencies in the particular oblique viewing directions used in these two cases. The surface orientation classification results for cylinder surfaces closely parallel the results found in the elliptical case.

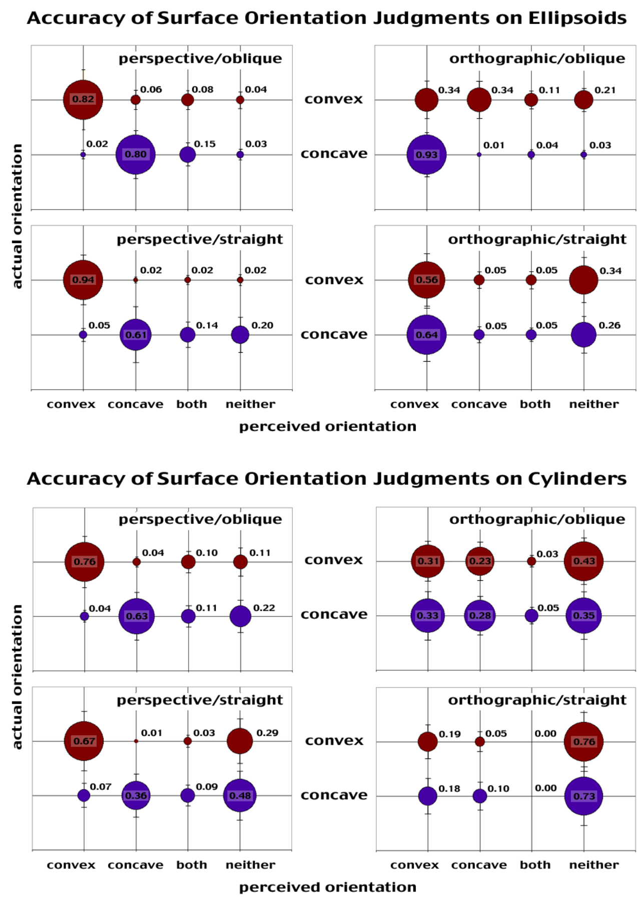

Figure 12: Rates of correct classification of surface orientation, for ellipsoids and cylinders, using results pooled over all texture types.

4.3 Discussion

The most interesting results from the first experiment are the findings that:

1) Shape classification accuracy rates are highest, in the general case, when surfaces are textured with the bi-directional pattern that follows the two principal directions (2dir). Looking at subsets of the results in detail we found no case in which shape classification was significantly worse with the 2dir texturing than with any other texturing condition. In most, 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 among those studied 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, if the surface was viewed from an oblique vantage point. In addition, shape classification accuracy was as good with the oblique/orthographic viewing as with the perspective/straight viewing, for many texture types. However, 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 interesting observations from the experiment are:

1) The diagonal textures caused particular problems for cylinder recognition. Results for ellipsoids was mixed: the 2dir45 texture worked well, but the 1dir45 pattern did not, presumably because it lacked the ability to convey curvature along more than one direction.

2) The noise and swirly textures performed best in the case of the ellipsoid surfaces, possibly because they exhibited the strongest gradients of texture spatial frequency due to distance (noticeable only under perspective projection).

3) The textures that followed just one of the principal directions frequently caused ellipsoids and saddles to be misperceived as cylinders, and sometimes (especially with orthographic projection or head-on viewing) caused cylinders to be misperceived as flat.

4) The three-directional texture seemed to work fairly well on the ellipsoid and saddle surfaces, but not as well on the cylinders. One possibility is that the presence of the diagonal component in the 3dir texture is interfering with shape perception in the cylinder case in the same way that it does in the 1dir45 case, which didn't include the pdir components. However a potentially complicating factor is that this texture, being drawn from a photograph, is slightly less rigidly regular then the other directional patterns, which were created artificially. Hence, it also could be the case that people were using a slightly different strategy when making classification judgments based on this texture. The conservative conclusion from the first experiment is therefore that adding information along more directions in addition to the principal direction does not appear to be 100% safe – there was not strong evidence of a positive effect, and there was possible evidence of a negative effect.

5. Experiment 2

5.1 Motivation

Although we were very pleased with the results we had obtained in the first experiment, there were several lingering questions from that experiment that we wanted to more directly address. The first concerned subtle doubts we had about the strength of our ability to reliably infer, from the results of that experiment, that there could be no inherent advantage in employing a texture that contained, in addition to indications of the first and second principal directions, a component oriented in a third direction oblique to those other two. The source of these doubts was twofold. First, in previous research seeking insight into the question of whether textures that followed two principal directions might inherently be better suited to showing shape than textures that followed only one, we had discovered that not all principal direction textures were equally effective for showing shape, and that subtle differences in texture characteristics could have significant impact on shape perception judgments. Second, we recognized that the particular texture pattern we had selected to exemplify the three directional case was a natural texture that we had aquired from the Brodatz album[18]. As such, it inevitably contained, and afforded the tolerance of, very subtle irregularities in the directions implied by the texture striations. Might it be possible that our finding that the 2dir pattern outperformed the 3dir pattern could be partly explained by this difference in texture character, leaving open the possibility that a different 3dir texture, which offered less affordance to the toleration of subtle irregularities, might yet have the potential to outperform the 2dir grid?

The second lingering question that motivated us to pursue a follow on experiment was the desire to look more closely and more rigorously at the nature and strength of the separate and combined contributions to shape understanding provided by each of the individual components in the multi-directional textures.

5.2 Design, Stimuli and Procedure

With these goals in mind, we defined a new set of texture patterns, built up incrementally from primary unidirectional components. Figure 13 shows six of the seven patterns used: 1dir, 1dir45, 2dir, 2dir45, 3dir1(=2dir+1dir45), and 3dir2(=2dir45+1dir). Pattern 1dir90 is not shown. Obviously this is not a complete enumeration of all possible combinations of uni-directional components, but because of practical considerations we needed to keep the number of different textures to a reasonable size, and we felt that this subset of patterns was sufficiently representative for the purposes of our investigations.

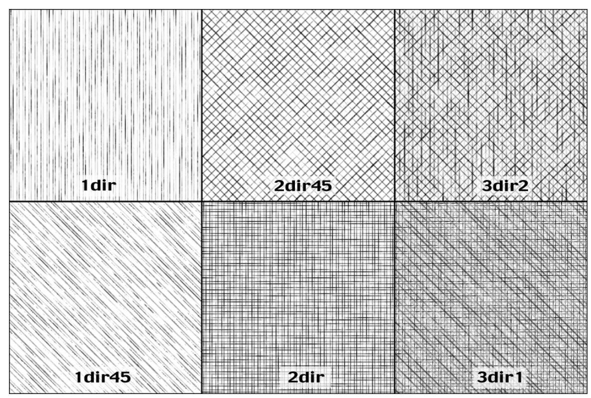

Figure 13: Six of the texture patterns used in our second experiment.

We used the same method of stimulus preparation we in our second experiment as in our first, except that this time we took exceptional pains to obtain synthesis results that were as perfectly regular and free of artifacts as possible. For reasons that are still not completely clear, we found that attaining good synthesis results for the 3dir patterns proved far more difficult than expected. Figure 14 shows the 2dir1, 3dir1, 2dir45, and 3dir2 patterns synthesized over the convex and concave elliptical and cylindrical surface patches, as well as the saddle and flat surfaces. We restricted our investigations in the follow-on study to the conditions of perspective projection and oblique viewing.

We recruited eight participants for our second study, two of whom had participated in the first study. The others were naïve to the research goals of the experiment. Participants viewed a total of 8x4x6=192 images: 8 texture types x 4 shape types x (2 orientations and 4 image rotations) for the ellipsoids and cylinders, or (1 orientation and 2 image ro tations) for the saddles and planes. Participants made the same shape classification and surface orientation judgments as before, though we only considered the shape classification responses in our analysis.

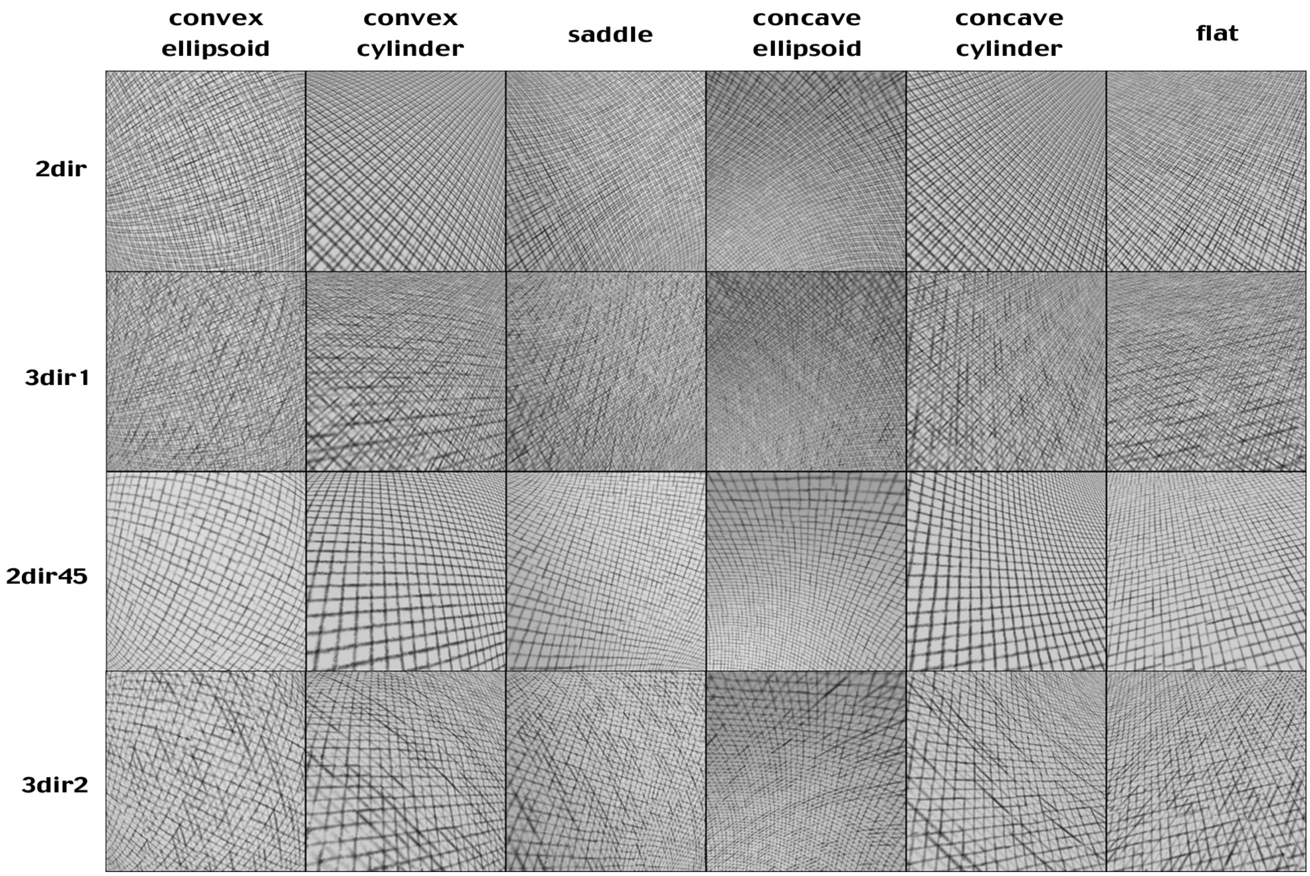

Figure 14: Example surface stimuli from our second experiment.

5.3 Findings

We carried out a 3-way ANOVA analysis (subject x texture_type x shape_type) on the combined results across all image rotations. We found significant main effects of subject ($F_{7,952}$=3.491, p<0.01), texture_type ($F_{6,42}$=36.94, p<<0.001) and shape_type ($F_{2,14}$=14.77, p<0.001), and significant two-way interactions between subject x shape_type, and texture x shape_type. Using the same Tukey HSD analysis as before, we found that shape classification accuracy rates were significantly better for the saddle surfaces than for the ellipsoids and cylinders (p<0.05; coeffs = {-0.103 cyl, -0.0424 ell, 0.145 sad}). In subsequent analysis, running the ANOVA on subject x texture x orientation, we found a marginally significant main effect of surface orientation on correct shape judgment in the case of the ellipsoids ($F_{1,7}$=5.072, p=0.05902), where the convex stimuli facilitated accurate shape categorization judgments to a greater extent than the concave stimuli (p<<0.001) but we did not find a significant effect of surface orientation in the case of the cylinders ($F_{1,7}$=0.9547, p=0.3611). The charts in figures 15-17 show details of our findings, using similar presentation formats as earlier.

The main significant, unexpected and surprising finding that we got from our second experiment is that performance was no longer best under the 2dir pattern condition in all cases. For correct recognition of cylinders, the 2dir texture was significantly outperformed by the 1dir, 3dir1 and 3dir2 patterns. Shape classification accuracy was also as good or better with the 3dir patterns in most other cases.

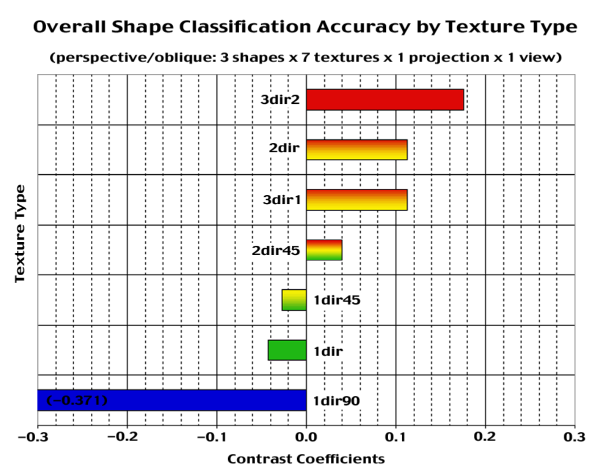

Figure 15: Results from the Tukey HSD analysis on the effects of texture type on shape classification accuracy in experiment, pooling results over all shape categories.

5.4 Discussion

There are several possible explanations for these startling results. One is that there were some subtle differences in the 2dir patterns used in first and second experiments, and the first principal direction line features appeared to be less prominent in some of the second set of images than in the first. Loss of this important signal in critical portions of the images might explain some of the observed effect.

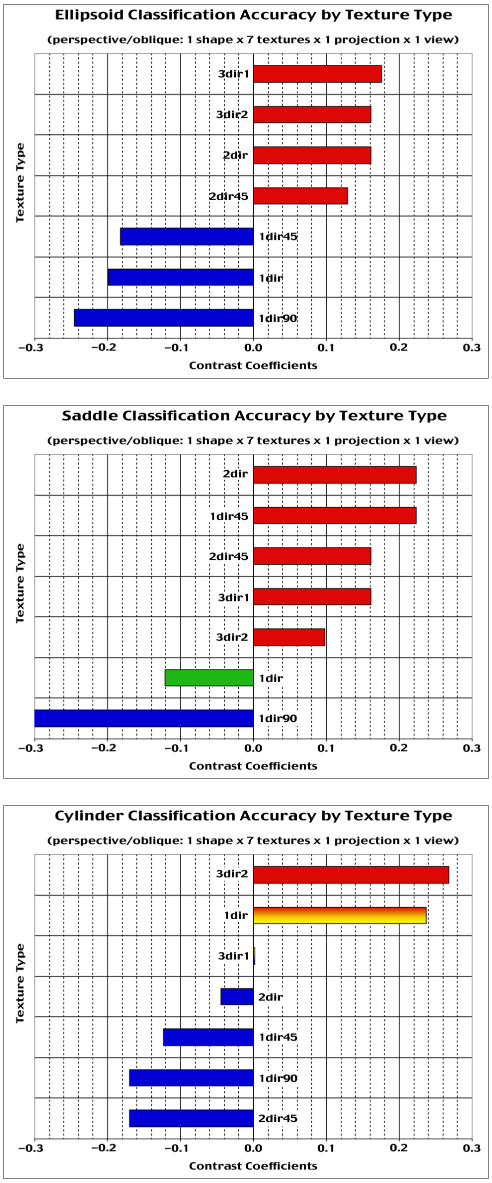

Figure 16: Results from the Tukey HSD analysis on the effects of texture type on shape classification accuracy in experiment 2.

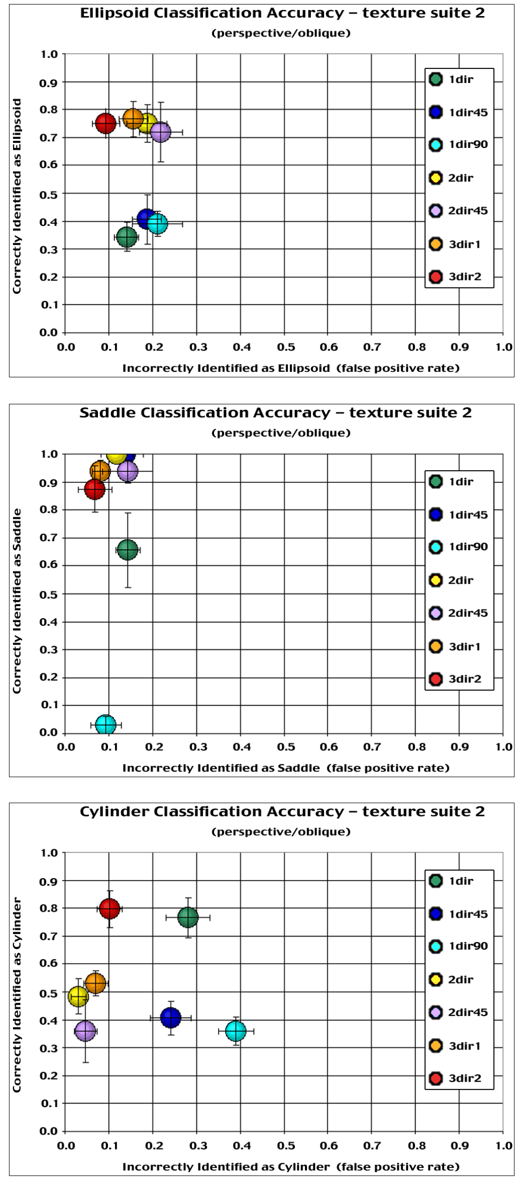

Figure 17: Charts showing shape classification accuracy results.

However, after the initial shock of the discovery wore off, a more intriguing possibilility occurred to us. Perhaps we had been too quick to put our faith in the inherent superiority of the principal direction grid pattern. It could be possible that the 1dir90 component in and of itself contributes remarkably little to shape perception — so little that replacing it with other balanced components, also capable of conveying distances along first principal direction, might enable showing shape as well or better.

This raises the importance of the issue of more rigorously determining how our visual system might go about constructing an inference of 3D shape from the information carried by the image of the flow of a texture pattern over the surface. Many researchers, such as Stevens [14] and Mamassian and Landy [15] have speculated that the visual system has a prior bias to interpret lines on a surface as if they were following the principal directions. But, this is a fairly strong assumption, and one that is not easily explained by any ecological theories of perception. It seems reasonable to infer that veridical shape perception might be optimal in the presence of a pattern containing indications of the intrinsic surface curvature in the first principal direction because images from a generic viewpoint of lines in any other direction would be less curved and hence give rise to a perception of the surface as being more flat. But this logic does not necessarily point to an important role for the second principal direction in shape perception, aside from the non-generic case in which the local shape has zero Gaussian curvature and the indication of the straight lines in the 2^{nd} princpal direction becomes a strong indication of the surface flatness.

An alternative possibility is that our interpretation of a doubly curved surface is indicated by the presence of certain critical skeletal shape features, such as the presence of pairs of curved lines that might plausibly lie in the surface. An indication of two non-intersecting lines that curve towards each other to enclose a convex region may be sufficient to create an impression of a local elliptical patch. An indication of two non-intersecting lines that curve away from each other may be sufficient to create an impression of a saddle-shaped region. If only one direction of curvature is implied, that may produce an impression of a cylinder.

6 Conclusions

In previous studies, reported elsewhere, we had found that observers' judgments of local surface orientation, made with a surface attitude probe, were most accurate under conditions of anisotropic texturing when the direction(s) of the texture 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.

In the current studies, we adopted an alternative metric for evaluating texture's impact on shape perception: comparing the relative accuracy with which observers were able to identify the intrinsic shape type (elliptical, cylindrical, hyperbolic or flat) and surface orientation (convex, concave, both, or neither), in close-up views of analytically-defined surface patches under a variety of different principal direction texture pattern conditions, plus two control texture conditions, under both perspective and orthographic pro jection and from both a head-on and an oblique viewpoint.

The most significant result of our present studies was to confirm the hypothesis that accurate shape perception is facilitated to a statistically significantly greater extent by some principal direction texture patterns than by others. In our first study we found that, under conditions of perspective projection for both viewpoint conditions, correct identification of both the shape type and surface orientation of the stimuli were best facilitated when the textured quadric patch was shown with a pattern that contained indications of both the first and second principal directions. In our second study, we found indications that surface shape classification judgments may hinge on the presence or absence of critical 'shape indicators' – lines, implied or explicit, that 'sketch out' the essential structural features of a form and which appear to be preferentially interpreted in a manner consistent with what would be the case if they followed a line of curvature on a convex form.

We believe that adopting an approach to evaluating shape perception that considers the understanding of intrinsic surface shape separately from extrinsic surface orientation, as has been previously suggested by Koenderink [19], can be valuable in enabling alternative possibilities for gaining deeper insight into the shape-from-texture process. Our finding that observers retained the ability, under many of the tested principal direction texture conditions, to make correct shape classification judgments in the absence of perspective projection offers a counterpoint to previous assertions that texture cannot carry shape information in the absence of perspective projection.

These are among the earliest studies that have sought to systematically investigate the effects of oriented texture components on shape perception for doubly curved surfacess. They have provided important insights but have also raised intriguing questions. Much work clearly re mains to be done before the question of how best to define a surface texture to facilitate accurate shape perception can be definitively answered.

Acknowledgments

This work was supported by NSF (ACI-9875368). We are grateful to the Statistical Consulting Clinic, which is supported by funding from the Minnesota Agricultural Experiment Station, for initial 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, Scholar-ship and Artistry. Research space was provided by the Digital Technology Center, which was funded by an initiative from the legislature of the state of Minnesota. We are indebted to our observers for their dedicated, tireless and conscientious efforts, and to the anonymous reviewers whose suggestions greatly improved this manuscript.

References

[1] V. Interrante and S. Kim, "Investigating the effect of texture orientation on shape perception," *Human Vision and Electronic Imaging VI*, SPIE 4299, 330-339, 2001.

[2] V. Interrante, S. Kim and H. Hagh-Shenas, "Conveying 3D shape with texture: recent advances and experimental findings," *Human Vision and Electronic Imaging VII*, SPIE 4662, 197-206, 2002.

[3] S. Kim, H. Hagh-Shenas and V. Interrante, "Showing shape with texture: two directions seem better than one," *Human Vision and Electronic Imaging VIII*, SPIE 5007, 332-339, 2003.

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

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

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

[7] B.G. Cumming, E.B. Johnston and A.J. Parker, "Effects of different texture cues on curved surfaces viewed stereoscopically", *Vision Research*, 33(5/6), 827-838, 1993.

[8] G. Gorla, V. Interrante and G. Sapiro, "Texture synthesis for 3D shape representation," *IEEE Transactions on Visualization and Computer Graphics*, 9(4), Oct-Dec 2003, 512-524.

[9] D. Knill, "Contour into texture: information content of surface contours and texture flow," *Journal of the Optical Society of America*, A, (1), January 2001, 12-35

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

[11] J.T. Todd and F.D. Reichel, "Visual perception of smoothly curved surfaces from double-projected contour patterns," *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 665–674, 1990.

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

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

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

[15] P. Mamassian and M. Landy, "Observer biases in the 3d interpre tation of line drawings," *Vision Research*, 38(18), 2817–2832, 1998.

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

[17] J.J. Koenderink, A. Doorn and A. Kappers, "Surface perception in pictures", *Perception*, **52**, 487–496, 1992..

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

[19] J.J. Koenderink, Solid Shape, MIT Press, 1990.