Visualizing intersecting surfaces with nested-surface techniques

Chris Weigle Russell M. Taylor, II

Department of Computer Science
University of North Carolina at Chapel Hill

Abstract

The need to compare surfaces occurs frequently in many scientific pursuits. Local comparisons can be quite difficult when the surfaces are placed side-by-side. But the two surfaces may intersect in complex ways if overlaid. Techniques have been previously developed for visualizing nested surfaces. The most successful of these apply opacity-modulated texture to the outer surface to better convey its shape while still limiting its occlusion of the inner surface. Although these techniques enable a better understanding the shape of two surfaces than does uniform opacity, they do not immediately lend themselves to intersecting surfaces. In this paper, we describe a simple geometric partitioning of intersecting surfaces that allows the application of existing nested surface techniques, such as opacity modulation, to a broader range of visualization problems. We also investigate how adding cast shadows to opacity-modulated textures affects the understanding of inter-surface distance and local shape.

Key words: Perception, non-photorealistic rendering, user study, scientific visualization, medical visualization

1 Introduction

In many scientific pursuits it is desirable to compare local and global shape between a pair or more of surfaces to understand how (and why) they differ. These surfaces may come from data acquisition or modeling and simulation. For instance, a physicist may wish to determine how well his specimen model and atomic-force microscope (AFM) simulation explain acquired AFM scans, or a doctor may wish to compare tumor images before and after a session of chemotherapy. As such, we are motivated to develop techniques that allow a human observer to clearly see relationships between two intersecting surfaces. While techniques exist that allow for the display of nested surfaces, intersecting surfaces have not been given much attention in the literature.

We are particularly interested in visualization techniques that enable scientists to perform multiple comparisons between surfaces, where the comparisons are intimately related to their research questions. A common difficulty in introducing new visualization techniques to a scientists workflow is convincing the scientist that investing time in a new visualization technique will yield better or faster understanding of data than already possible with display methods to which they have grown accustomed. To this end, we collaborate with scientists to identify measurable performance tasks that match well with their research questions. We then conduct evaluation studies to predict how well our new visualization techniques convey shape information scientists can use to explore their data.

Many techniques for nested surfaces solve the occlusion problem by displaying the outer layers with modulated opacity [15, 25] This allows portions of the inner surfaces to be visible through portions of the outer surfaces. Intersecting surfaces pose a more difficult problem, since no one intersecting surface can be said to be inner- or outer- most. We describe a simple repartitioning of the surface geometries at the intersections that allows existing nested-surface techniques to be applied to intersecting surfaces. We describe the results of a user-study evaluation of our glyph-like variant of one of the more successful nested-surface techniques (Interrante's curvature-directed strokes) for intersecting-surface display, and investigate whether adding cast shadows improves perception of inter-surface distance without significantly impacting local surface orientation estimation.

2 Displaying a Single Surface

Many published perceptual studies explore how best to present the shape of a single surface. Several studies in visual perception document the ability to perceive shape from shading [1, 4, 6, 7, 12, 24]. Specifically, the human visual system interprets shape with a built-in bias toward scenes that are diffusely lit from overhead [12] - though sufficient cues from other perceptual channels can influence this bias. Perception research suggests a variety of other visual cues that also illicit the perception of shape, such as texture [6], specular highlights [1], shadows [10], and object boundaries [24].

The most common illumination model used in computer graphics, the empirical Phong lighting model, may convey shape cues in a manner similar to certain real objects under natural illumination [16]. The Phong lighting model approximates both diffuse and specular lighting according to Lambert's Law and Snell's Law, respectively [23]. Under the appropriate conditions, Phong illumination has been shown to convey shape and depth.

Much research into the perception of shadows has considered attached shadows [10, 2], and how they are perceived similarly to occlusion contours. Cavanagh and Leclerc report that cast shadows help disambiguate light direction and help place objects in space, but do not appear to strongly suggest shape [3]. Mamassian et al. report that even when available visual information suggests otherwise, shadows in motion are interpreted as if the light source were fixed [21].

Texture has long been known to be an excellent shape cue. Gibson showed that the slant of a textured plane extending to the horizon could be perceived, if underestimated [12]. Cumming et al. described the three components of a uniformly textured surface, *compression*, *density*, and *perspective* [5]. Of these, texture *compression* – the apparent change in the texture due to the difference between the surface normal and view direction – has been shown to be the most significant for surface shape perception under stereo viewing.

A number of studies have found that textures can be constructed specifically to enhance shape and depth perception better than an arbitrarily constructed texture. Li and Zaidi find that (sinusoidal) surface shape is best perceived when noise-like textures project significant energy along the first principal curvature direction [20] – the direction of highest local curvature. Interrante et al. found that brush strokes laid along the first principal curvature direction, through well spaced points on the surface, also convey surface shape [15]. Kim et al. showed that using both principal curvature directions could convey shape better than either one alone [17].

Recently, some have cast doubt on how important the first principal curvature direction is in the perception of shape from texture [27, 26]. The argument made by both deals with the loss of the shape cues (from texture compression) as the projection of the principal curvature direction comes to coincide with the view direction. Sweet and Ware point out that using the direction orthogonal to the first principal curvature direction produces much stronger perception of shape in these cases [26]. But this is exactly the second principal curvature direction – the direction of least local curvature. As noted earlier, it has been suggested that using both principal directions together provides better shape perception than either one alone [17].

Sweet and Ware describe a highly effective and inexpensive two-direction texture hash [26]. The texture is computed as the intersection of two sets of regularly spaced infinite planes with the surface, each set of planes orthogonal to the other and to the view direction. This produces a grid on the surface that provides texture compression cues and can be used to judge distances along the surface. The texture pattern's dependence on the view direction means it is not constant through changes in view, however, and as such would not be appropriate for interactive exploratory visualization.

3 Displaying Nested Surfaces

Several techniques have been developed to allow the visualization of nested surfaces. All of these techniques can be said to fit in the broad category of non-photorealistic rendering (NPR).

Uniform translucency confounds shape perception away from object silhouettes. Diepstraten et al. describe a technique for view-dependent transparency, which aims to automatically produce translucent surfaces similar to technical illustrations [8]. In a later work, the same authors describe techniques for automatically producing breakaway and cutaway illustrations of nested surfaces [9]. These illustrations remove portions of geometry that occlude the interior surfaces instead of rendering them translucently.

Many techniques render an opaque interior surface surrounded by textured, translucent surrounding surfaces [15, 25]. The texture patterns modulate local opacity, providing better illumination and texture cues to enhance exterior surface shape perception as compared to uniform translucency. Interrante placed first principal curvature direction strokes along outer surfaces [15]. Rheingans retiled surfaces so that uniform circle or hexagon textures could be applied around vertices [25]. Interrante provides an excellent summary of relevant perceptual issues for such visualizations [15].

4 Displaying Intersecting Surfaces

Nested surfaces offer a single occluding surface to which one can apply a rendering device that allows the interior surface to be partially visible. Intersecting surfaces do not. Which surface occludes the other is view-dependent. For most interesting surfaces, it is likely that each surface occludes the other in different regions of a single view. It is still possible to use nested surface techniques, however. One only has to determine what parts of the surface are interior or exterior.

We propose using the intersections to partition the source surfaces into sections. Each section can then be labeled, and nested surface techniques applied according to label. We propose that each section have two labels. The first label identifies to which source surface it belongs.

The second label identifies if it belongs to the group of sections *interior* to any other surface or the group *exterior* to all other surfaces.

Determining if sections of the surface are *interior* or *exterior* can be accomplished in many ways. Exact geometric solutions are available though, for rendering alone, an image-space technique is sufficient. Ray casting is probably the simplest such solution, conceptually and in terms of implementation. A variety of z-buffer techniques exist [11, 8, 13, 22]. We used a hybrid of order-independent transparency and parity testing which operates like a ray-caster. The algorithm operates completely on the graphics hardware and finds the intersections implicitly.

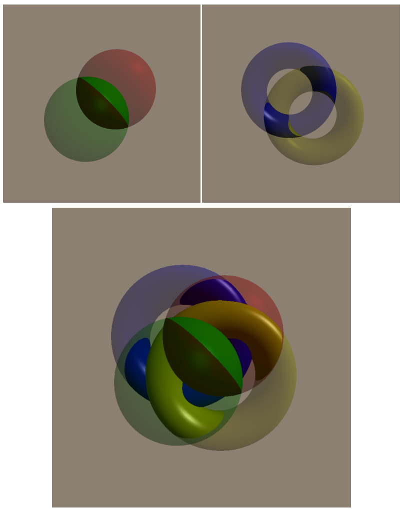

Figure 1: These simple objects intersect eachother. By partitioning the geometry at the intersection, we can render the interior and exterior differently. Here the interior is opaque and the exterior is translucent.

Once labeled, we can apply perceptual tags to the surface sections for rendering. Figure 1 presents some simple intersecting objects displayed with the *exteriors* translucent and the *interiors* opaque. For evaluation in this paper, we chose to use color to identify the different source surfaces and texture and opacity to distinguish between interior and exterior. *Interior* sections are rendered opaque and textured with a grid. *Exterior* sections are rendered with an opacity-modulating texture composed of glyphs. The glyphs are crosses oriented along principal curvature; each glyph has a long arm along the first principal curvature direction and a short arm along the second (see Figure 2).

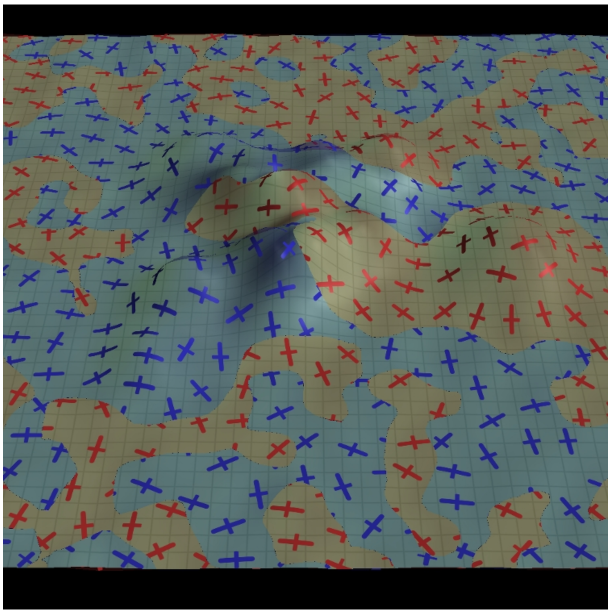

Figure 2: This is an example of a glyph visualization from our user study. Note that the interior is a different color from either exterior; we explain why we changed the color later.

5 Design

In this section, we discuss the design of the user studies performed to evaluate how well our principal-direction glyphs convey the shape of intersecting surfaces.

5.1 Tasks

The shape tasks used in the user studies were chosen in consultation with our collaborators in the sciences. From our standpoint, we had three goals:

- to satisfy the scientists that the study tasks related to their research questions,
- to use shape tasks commonly employed in the literature (adapted for two simultaneous surfaces),
- and to facilitate quick, simple judgments from the participants.

For this paper, we interviewed a physicist and a surgeon. The physicist is interested in comparing real atomic-force microscope (AFM) scans with simulations. The physicist's goal is to validate his model of the structures of specimens. The AFM collects 2D height data by scanning a crystal tip across the surface of the specimen, effectively computing a dilation of the true specimen by the tip shape. Understanding the differences between the real and simulated scans can help determine what changes need to be made to the model. For instance, if peak heights are not significantly different between the two surfaces, but the slopes are different, the most likely explanation is that the tip model is the wrong size or shape. The surgeon investigates methods of automatic image segmentation of tumors from magnetic-resonance imaging (MRI). Understanding the differences between automatic segmentation and expert-human segmentations can help determine how to tune the automatic algorithms. For instance, if the automatically-generated surface consistently (multiple occurrences per single data set) produces a particular feature differently than the expert-human, it is likely that the algorithm needs tuning.

Generally, these scientists need to understand shape and inter-surface distance. Moreover, the scientists should be satisfied that the user studies predict how well our visualizations convey both. For this reason, we will perform studies to separately evaluate each. Langer and Bülthoff lists the most commonly used shape tasks and weighs their advantages and disadvantages [19]. We chose to ask participants to compare properties at a pair of points. Since displays represent two surfaces, the difference between the two surfaces should be evaluated at each point. This requires making a forced choice between the two points instead of manipulating some measurement widget (like the orientation probe developed by [18]) once for each of the two surfaces represented by the display. Participants can therefore perform many trials in a relatively short amount of time. However, it does mean that we will have simple binary (right or wrong) responses to evaluate instead of a measurement error. We describe each task in more detail below in the Experiment sections.

5.2 Data and Presentation

Each pair of surfaces is computed on a 100x100 grid. Inter-surface distances are computed at ten times that resolution. Each surface is comprised of eight Gaussian bumps. The Gaussians have random centers, maxima, and sigmas, and may overlap. Noise is also included in each surface; the noise is significantly smaller in scale than the Gaussians. Figure 3 shows several examples.

Renderings of the surfaces are pre-computed as if viewed from 53 cm away from 45 degrees above the plane of the grid (around the x-axis). The initial renderings are over-sampled, filtered, and re-sampled at the final display resolution to reduce aliasing artifacts. As rendered, the surfaces extend horizontally just beyond the viewport, which is 22 cm wide. The viewport is 24 degrees square.

Illumination is from an infinite point source with a vector $(1, .5, 1)$, effectively over the viewer's right shoulder. Per-pixel Phong illumination is employed.

When surfaces are first presented, they rock five degrees to either side of center around the y-axis. The motion is that of a torsion pendulum, slowing to rest before direction changes. Participants may also repeat the rocking animation as they desire during trials.

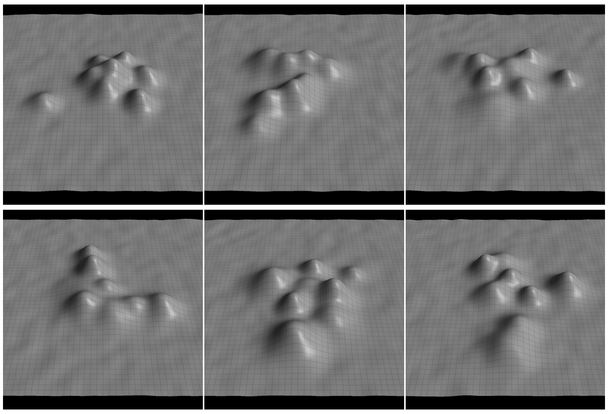

Figure 3: These are examples of the random surfaces created for the evaluaton.

5.3 Conditions

The first visualization condition is a single surface displayed with color mapped to the signed distance to the closest point on the second surface (Figure 4). A double-ended saturation scale, blue through gray to red, classifies inside, outside, and touching (hue) while also encoding distance (saturation). The choice of blue is a trade-off. The human visual system has reduced sensitivity to blue compared to red or green, but using this color include participants with red-green colorblindness (the most common form). Color mapping is a frequently-used visualization technique for conveying scalar parameters on surfaces – a task for which it is well suited, when used appropriately [14, 28]. Applying a color mapping to a surface can be quite effective for classification or metric estimation, but is not a good perceptual channel for carrying *alternate* shape information. Unfortunately, it is also often used to express the differences between two surfaces – a task for which it is of limited use. Color can be used to classify whether a hidden surface is *inside* or *outside* the displayed surface. Color can be used to encode the distance between points on the displayed and hidden surfaces. But neither of these enables true shape perception.

The second visualization condition is the textured glyph visualization we have described previously (Figure 5). Red and blue are used to distinguish the *exterior* of the two surfaces. The *interior* surfaces are colored light gray instead of according to the source surface as we have suggested before. We do this to conform to the third condition.

The third visualization condition adds cast shadows to

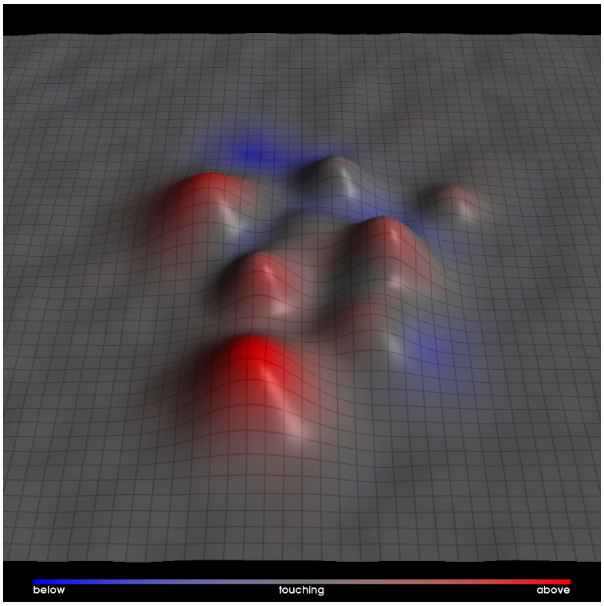

Figure 4: This is an example of the color condition. The surfaces include the top middle example surface from Figure 3.

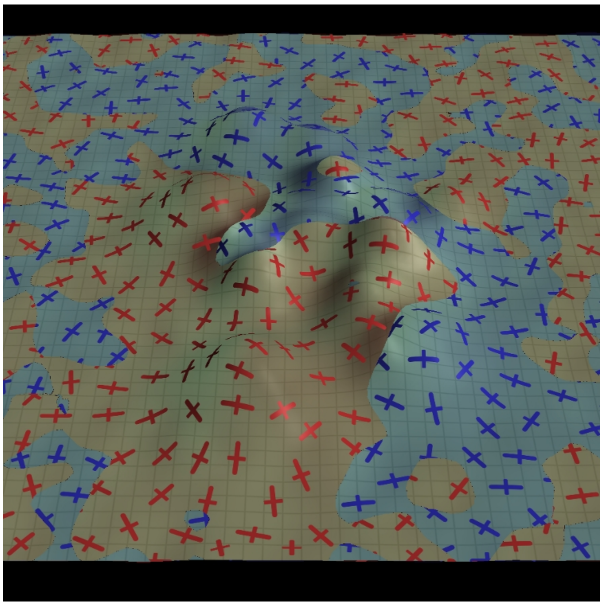

Figure 5: This is an example of the glyph-without-shadows condition. The surfaces include the top middle example surface from Figure 3.

the textured glyphs (Figure 6). Red and blue distinguish the *exterior* portions of the two surfaces. The light gray *interior* provides higher contrast for the shadows cast by the glyphs. These cast shadows should enhance the ability to perceive inter-surface distances, a metric of interest to our collaborating scientists, as compared to the unshadowed glyphs.

6 Evaluation

In this section, we describe the two user studies and analyze the results of each. The two experiments were run in parallel, so insights from one experiment could not be applied to improve the other.

6.1 Experiment 1 - Distance

Six graduate and undergraduate students participated in this experiment. They were each asked to view sixty (60) unique, pre-generated, random surface pairs per visualization condition, for a total of 180 trials. In each trial, two points were marked with overlaid circles (1.5 cm in diameter), one cyan and solid, one yellow and dashed. The participants were asked to estimate how close the two surfaces approached each other within each circle, and report the circle containing the closest approach. The circles were pre-determined to guarantee that the per pair differences in closest approach were uniformly distributed among 0.0 through 5.5 grid units in 0.5 unit intervals. Trials were randomly ordered for each participant. Figure 7 is an example of a trial as presented to the participant.

We expected that the accuracy of participant responses should depend on the magnitude of the differences between the distances compared and on the visualization condition. We expected that the cast shadows should enhance the perception of separation between surfaces. We expected that participants should be able to estimate distances accurately directly from the color.

6.2 Results of Experiment 1

ANOVA analysis of main effects finds significance for distance difference ($p << .001$), visualization condition ($p < .01$), participant ($p < .05$), and response time ($p << .001$). Figure 8 shows the accuracies predicted from a logistic regression over distance difference and visualization condition. From the fit, we see clearly that the two glyph techniques are not separable, but both enabled more accurate responses than color.

Responses to questionnaires show that the average participant found color and glyphs with cast shadows to show the distances with equal clarity and with greater clarity than glyphs alone. However, the average participant preferred glyphs with shadows over the other two techniques.

6.3 Experiment 2 - Orientation

Seven graduate and undergraduate students participated in this experiment. Each responded to 180 trials, simi-

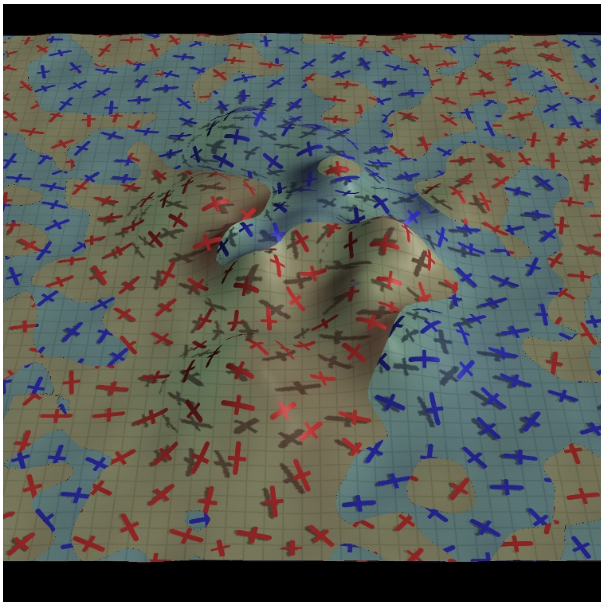

Figure 6: *This is an example of the glyph-with-shadows condition. The surfaces include the top middle example surface from Figure 3.*

lar to Experiment 1. Circles marked the points of interest, also as in Experiment 1. The participants were asked to estimate the orientation of each surface within each circle and report the circle containing the smallest orientation difference. The difference-angles between orientation differences were uniformly distributed among 0.0 through 44.0 degrees in 4.0 degree intervals. [**?**] note that orientation differences between points on a single smooth surface can be detected at different thresholds depending on the separation between the points. The just-noticeable-difference in angle could be as little as 8 degrees for neighboring points and as much as 22 degrees for widely separated locations [**?**]. Again, the circles were pre-determined and the trials ordered randomly. Figure 9 is an example of a trial as presented to the participant.

We expected that the accuracy of participant responses should depend on the magnitude of the between-marker difference of the between-surface orientation differences and on the visualization condition. We expected that the cast shadows should enhance the perception of shape, especially on the *interior* surface. We expected that participants should be able to estimate orientation differences accurately by detecting gradient magnitudes in the vicinity of the selected regions. The distance between the two surfaces should be changing wherever the two surfaces are not parallel, and so the rate of color change indicates the difference in orientation between the two sur-

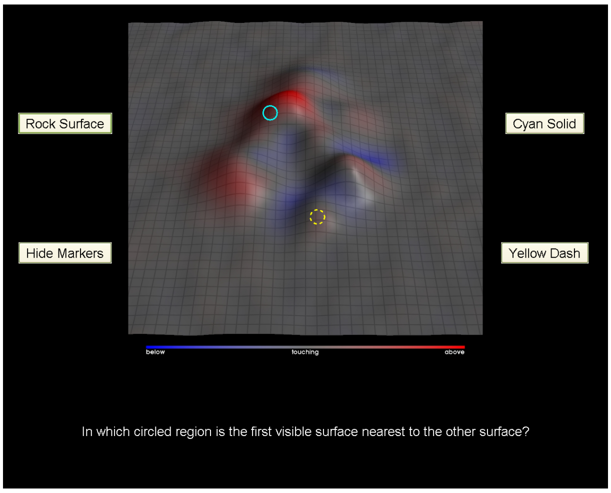

Figure 7: *This is an example trial from the distance task.*

faces. The systematic bias in estimating orientation often reported in the literature should cancel out in the glyph conditions, as both circled regions contain the same two perceptual techniques (one each of a glyph-textured surface and a grid-textured surface).

6.4 Results of Experiment 2

ANOVA analysis of main effects find significance for angle difference ($p << .001$), visualization condition ($p < .01$), and response time ($p < .01$). Figure 10 shows the predicted accuracies from a logistic regression over angle difference and visualization condition. From the fit, we can again see that the glyph techniques enable more accurate responses than color, but can not be separated from each other. Notice that for larger angle differences the color technique begins to approach the two glyph techniques again. In these cases the color gradients are large in magnitude and compressed over a small space.

Responses to questionnaires show that the average participant found glyphs with shadows marginally clearer than color, and either clearer than glyphs alone. No preference was reported between color or glyphs with shadows, but either were preferred over glyphs alone.

6.5 Discussion

Overall, the glyphs techniques appear to provide for improved shape estimation over mapping inter-surface distance to color. More interesting, the glyph techniques, which show portions of both surface geometries, also appear to convey inter-surface distance better than directly mapping that distance to color on a single surface. It must be said that the color map chosen is not the optimal map for displaying a signed, continuous metric on a smooth

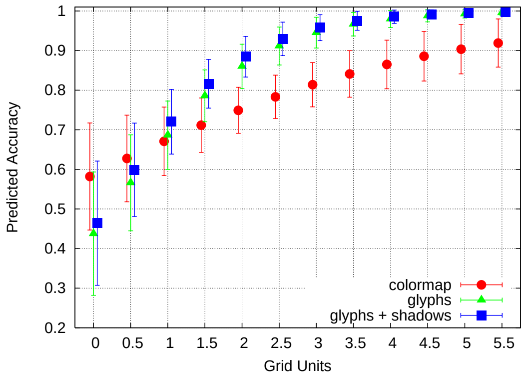

Figure 8: The figure shows predicted accuracy from the fitted model for estimating differences in inter-surface distance. The model accounts for the difference in distances and the visualization technique.

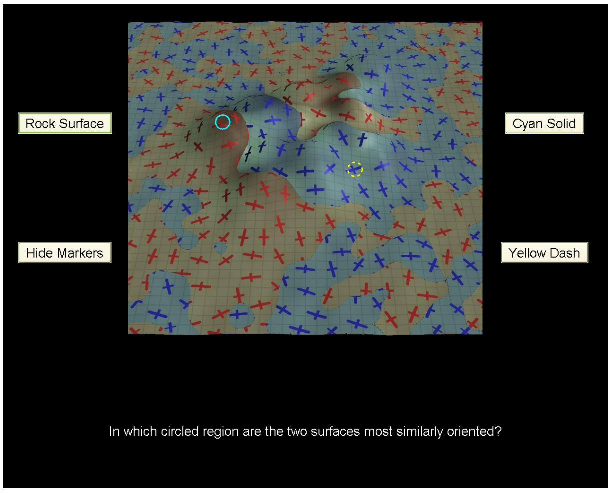

Figure 9: This is an example trial from the orientation task.

surface. It is also not the optimal map for perceiving gradients of that metric. However, it is a likely naive choice.

There are a number of reasons why adding cast shadows to the glyph technique does not have a measurable effect. One reason is that shadows may make no difference. Another reason we may have failed to find a difference is statistical power – many more participants may be required to separate the two conditions. It is possible that the rocking animation reveals too much of the shape for differences in presentation of the glyph technique to have any effect. The static images of the two glyph techniques certainly seem to suggest that shadows do affect perception (and they should help ground the location of the glyphs floating in space). Participants also spent a significant amount of time rocking some surfaces before responding.

7 Summary

We have described a simple partitioning of intersecting surfaces that enables existing nested surface techniques to be applied to a new visualization problem. By partitioning geometry at the surface intersections, sections of the intersecting surfaces can be categorized into *interior* and *exterior*. The sections can then be displayed as nested surfaces.

We performed user studies to show that a glyph-based variant of Interrante's curvature-directed strokes can be applied to convey local shape for intersecting surfaces. We were unable to find any significant difference between using glyphs with or without shadows. This is likely a side-effect of rocking the surfaces; the animation reveals much of the surface shape in the unshadowed case. Participants reported preferring the glyphs with shadows and felt they perceived shape better as well.

An important open question is whether it is still possible to perceive the global shape of one complete surface in either of the glyph visualizations, or if the full visualization is too complex and confounds such a task. We have undertaken to study this question, but have as yet no experimental results to report.

References

[1] A. Blake and H. Bulthoff. Does the brain know the physics of specular reflection? *Nature*, 343:165–168, 1990.

[2] U. Castiello. Implicit processing of shadows. *Vision Research*, 41:2305–2309, 2001.

[3] P. Cavanagh and Y. Leclerc. Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance*, 15(1):3–27, 1989.

[4] Chris G. Christou and Jan J. Koenderink. Light source dependence in shape from shading. *Vision research*, 37(11):1441–1449, 1997.

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

[6] William Curran and Alan Johnston. The effect of illuminant position on perceived curvature. *Vision research*, 36(10):1399–1410, 1996.

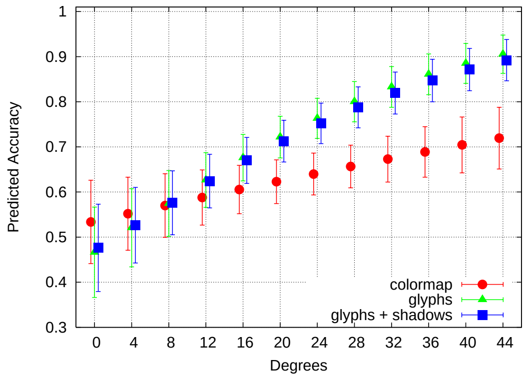

Figure 10: The figure shows predicted accuracy from the fitted model for estimating angle differences between set of local surface orientation. The model accounts for the difference in angle between pairs and the visualization technique.

[7] Erik De Haan, Roderik G. F. Erens, and André J. Noest. Shape from shaded random surfaces. *Vision research*, 35(21):2985–3001, 1995.

[8] J. Diepstraten, D. Weiskopf, and T. Ertl. Transparency in interactive technical illustration. *Computer Graphics Forum*, 21(3), 2002.

[9] J. Diepstraten, D. Weiskopf, and T. Ertl. Interactive cutaway illustrations. *Computer Graphics Forum*, 22(3):523–532, 2003.

[10] R. Erens, A. Kappers, and J. Koenderink. Perception of local shape from shading. *Perception and Psychphysics*, 54:145–147, 1993.

[11] C. Everitt. Interactive order-independant transparency. Technical report, Nvidia, 2001.

[12] J. J. Gibson. *The perception of the visual world*. Houghton Mifflin, Boston, 1950.

[13] J. Goldfeather, S. Molnar, G. Turk, and H. Fuchs. Near real-time csg rendering using tree normalization and geometric pruning. *IEEE Computer Graphics and Applicaitons*, 9(3):20–27, 1989.

[14] Chistopher Healey. Choosing effective colours for data visualization. In *Proceedings Visualization '96*, pages 263–270. IEEE Computer Society, 1996.

[15] Victoria Interrante, Henry Fuchs, and Steven Pizer. Conveying the 3d shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117, 1997.

[16] Alan Johnston and William Curran. Investigating shape-from-shading illusions using solid objects. *Vision research*, 36(18):2827–2835, 1996.

[17] Sunghee Kim, Haleh Hagh-Shenas, and Victoria Interrante. Showing shape with texture: Two directions are better than one. In *Human Vision and Electronic Imaging VIII*, 2003.

[18] J. J. Koenderink and A. Van Doorn. Relief: pictoral and otherwise. *Image and Vision Computing*, 13(5):321–334, 1995.

[19] M. S. Langer and H. H. Bulthoff. Measuring visual shape using computer graphics phychophysics. In *Proceedings of Eurographics Workshop on Rendering Techniques*, pages 1–9. Springer, 2000.

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

[21] Pascal Mamassian and Daniel Kersten. Illumination, shading and the perception of local orientation. *Vision research*, 36(15):2351–2367, 1996.

[22] F. Nooruddin and G. Turk. Interior/exterior classification of polygonal models. In *Proceedings of IEEE Visualization '00*, pages 415–422, 2000.

[23] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18:311–317, 1975.

[24] V. Ramachandran. Perceiving shape from shading. *Scientific American*, 259:58–65, 1988.

[25] Penny Rheingans. Opacity-modulating triangular textures for irregular surfaces. In *Proceedings of Visualization '96*, pages 219–ff. IEEE Computer Society, 1996.

[26] G. Sweet and C. Ware. View direction, surface orientation and texture orientation for perception of surface shape. In *Proceedings of Graphics Interface 2004*, 2004.

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

[28] C. Ware. Color sequences for univariate maps: Theory, experiments and principles. *IEEE Computer Graphics and Applications*, pages 41–49, September 1988.