

A Grid with a View: Optimal Texturing for Perception of Layered Surface Shape

Alethea Bair and Donald House, *Member, IEEE*

Abstract— We present the results of two controlled studies comparing layered surface visualizations under various texture conditions. The task was to estimate surface normals, measured by accuracy of a hand-set surface normal probe. A single surface visualization was compared with the two-surfaces case under conditions of no texture and with projected grid textures. Variations in relative texture spacing on top and bottom surfaces were compared, as well as opacity of the top surface. Significant improvements are found for the textured cases over non-textured surfaces. Either larger or thinner top-surface textures, and lower top surface opacities are shown to give less bottom surface error. Top surface error appears to be highly resilient to changes in texture. Given the results we also present an example of how appropriate textures might be useful in volume visualization.

Index Terms— Perception, optimal visualization, texturing, layered surfaces.

1 INTRODUCTION

Layered surface visualization is an important but difficult perceptual problem. It has many uses, including visualizing tissue layers in medical visualization, geological layers, or isosurfaces in technical illustration. However, visualizing layered surfaces is difficult because the layers tend to visually interact. A simple example is opacity; a fully opaque top surface is clearly visible but makes seeing the bottom surface impossible, while a fully transparent top has the opposite effect. Some balance must therefore be found to achieve a visualization that optimally shows relevant characteristics such as overall structure or specific features. Surface texture is a key factor that has been shown to be useful.

To our knowledge, the only work that has studied stylistic aspects of layered surface texturing is a series of studies by the authors and C. Ware [2, 3, 11, 12, 13]. These studies showed, among other things, that bright surfaces and bigger textures on the top surface might be better for layered surfaces. We continue this research here using two controlled perceptual studies to investigate texture spacing and opacity variables.

2 BACKGROUND

Perceptual research has shown that humans use a variety of visual cues to extract 3D shape from an image. These include shading and occluding contours [24], motion [21], stereo, lines and texture. In line drawings, studies have shown that both adults [8] and infants [5] search for line junctions in an image, and combine these to derive a 3D shape. Work has been done to catalogue various line junctions and possible interpretations [4]. Some evidence suggests that line direction, when drawn as a surface contour is interpreted as being a geodesic, or line of curvature [20]. However, ample evidence exists that humans can perceive shape from other lines, such as contour patterns [28], using a parallelism assumption. Under the assumption of isotropic textures, visual shape cues include compression, size gradient, density gradient, compression gradient, and perspective convergence [25]. However, evidence exists that humans can derive information from anisotropic and inhomogeneous textures [29].

- Alethea Bair is with Texas A&M University, E-Mail: abair@viz.tamu.edu.
- Donald House is with Texas A&M University, E-Mail: house@viz.tamu.edu.

Manuscript received 31 March 2007; accepted 1 August 2007; posted online 27 October 2007. Published 14 September 2007.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

Considerable research has already been done to find optimal visualizations for single surfaces. Surface shape judgments are improved by occluding contours [23], an oblique viewing angle [26], and motion cues and stereo [31]. Numerous studies have all shown that applying certain types of texture to a surface is better than simply shading the surface for direction or curvature judgments [7, 19, 26, 27]. Interrante and Kim [15] showed that both principal-direction texture and uniform direction texture seem to be better than random and non-surface-dependent textures for showing shape of a surface. Later, Interrante *et al.* [14] showed that principal-direction textures seem better than uniform textures, although the results include a viewing angle from directly above, a direction from which a uniform texture gives no information from surface contours. Results from other studies suggest that two-directional textures outperform one-directional textures in many cases [18, 19, 26].

For layered surfaces with two layers, Interrante *et al.* [16] showed that including opaque texture elements on the top surface significantly improves the accuracy of shortest-distance judgments. Black and Rosenholtz [6] present a successful shape-from-texture algorithm for occluded textures, but do not look at simultaneous processing of two textured shapes.

Stylistic aspects of texture such as color, size, opacity, shape and randomness are extensively used in data visualization to encode variables. Research has also been done on the effects of these variables in 2D visualizations [1, 22, 30]. In related research, Interrante and Grosch [17] found that color contrasts were not sufficient to depict overlapping streamlines under non-stereo viewing. This is similar to results found in the studies by House *et al.* [3, 12, 13], in which hue was not a significant factor in layered surface tasks.

3 EXPERIMENT ONE

In our first experiment we sought to answer the following questions. How does layering of two surfaces affect the amount of shape information that can be shown for terrain surfaces? How do well-textured layered surfaces compare to un-textured ones? Finally, how does spacing of the textures affect the quality?

3.1 Experimental Setup

Figure 1 shows a typical surface used in our studies. The surfaces were height fields constructed from Gabor bumps [6] to give a terrain-like feel. The surfaces were displayed full screen at a distance of 85cm from the subject, on monitors 48 cm wide and 30 cm high, with a screen resolution of 3840 x 2400. The surface construction algorithm drew 100 bumps on each surface, randomizing position,

amplitude, orientation, cosine period and the Gaussian falloff parameter σ for each bump. The cosine period varied between 7.5-20% of the surface width, and σ varied from 40-70% of the period. This made features with about 1-6° of visual angle. Surfaces were displayed in stereo using a Wheatstone stereoscope setup [32] as shown in Fig. 1, and rocked to provide motion cues to depth. The camera positions and field of view were adjusted to make the leading edge of the top surface appear to have the same depth as the physical monitor. At these settings, the upper and lower surfaces would appear to have a vertical separation of about 5cm, with maximum height variations of around 4cm.

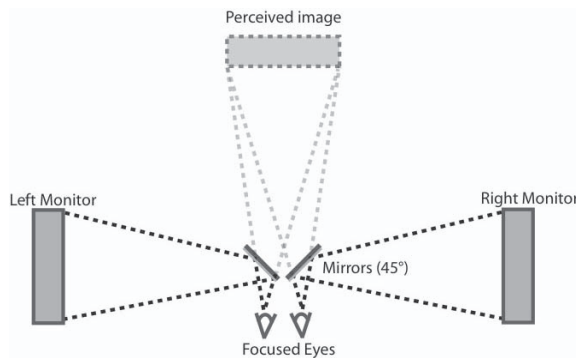


Fig. 1. Wheatstone stereoscope setup.

Subjects were given the task of manipulating a direction probe to align it with the perceived surface normal at a given point on either the top or bottom surface. Figure 2 shows the probe, which is of a new, unpublished design by Colin Ware. The probe consists of a fully three-dimensional cylindrical pole with an applied stripe texture, topped by a torus that is at right angles to the pole. The ellipse of the projected torus helps show the probe direction with minimal occlusion of both the cylinder and underlying surface. The torus is at the top of the cylinder because if it were at the bottom, occlusion with the surface would give additional shape information. Finally, the stripes mirror the torus direction as well as providing high-frequency texture to aid stereovision. The probe is rendered with the same lighting and shading as the surface. Thus, probe visual orientation cues come from the shape, shading and foreshortening of the texture stripes. This design makes the probe direction much clearer than the simple line probes used in previous research [15], and does not require a separate enlarged view of the probe [26]. Measured errors should be primarily due to errors in surface direction estimation rather than probe direction estimation, however, we have not done a controlled evaluation of this design verify this.




Fig. 2. Gabor surface and probe

Subjects were first trained to manipulate the probes while being shown the correct normal, and a color-coding system indicating their current angular error. Next, subjects were given a series of 'practice' probes in which they aligned the probe to the best of their ability, and

then only later were shown the correct normal and their angular error. Figure. 3 shows zoomed in views of the setup used for the training phase. During the actual experiment, subjects were not shown their errors. All training was done on a single surface to avoid bias in training.




Fig. 3. Examples of well and poorly aligned probes.

There were 14 subjects, (11 students and 3 faculty, 6 female and 8 male), from the Texas A&M Visualization Laboratory. Subjects had normal or corrected-to-normal vision, and were tested for stereovision on our equipment and color-blindness. All subjects had significant experience in visualization, and previous familiarity with the concept of surface normals. However, most were naive to the purpose of the experiment.

3.2 Fixed Texture Parameters

Following guidelines strongly suggested by previous research [3, 13, 26], textures were chosen to be uniform grids with high luminance. Although a previous study found that bottom surface textures should be subtler for a target-finding task [3], our preliminary tests showed that these subtle textures did not work well for the task of orienting normal probes. Presumably, subtle bottom surface textures are better for a task that simply requires recognizing the presence or absence of bumps because shading information is sufficient to distinguish a symmetric bump from a noisy background. Also, a subtle texture on the bottom reduces texture interference while still helping to distinguish the two surfaces. However, in the case of a probe-alignment task, subtle textures on the bottom did not prove adequate, so color, value and saturation were made equal on both top and bottom surfaces. Top surface texture backgrounds were given 50% opacity so that approximately the same pixel intensity contribution was made by each surface. The grid lines were made opaque to give high texture contrast between the lines and holes. Grid spacing was chosen to be smaller than the smallest surface feature size, but big enough to be easily recognized as a grid. The three spacings chosen had 100, 140 or 180 grid lines along a surface edge, which roughly correspond to a grid spacing of 0.6, 0.4 and 0.3 cm seen at 85cm. The top surface texture was rotated 45° relative to the bottom texture to minimize Moire patterns from near-parallel lines. Finally, hues were chosen giving a red top surface, blue bottom surface and green probe. This was solely for aesthetic reasons, since House *et al.* [13] showed that choice of surface hue is unlikely to affect surface shape perception. Color cues probably do not help significantly when the much stronger cues of shading, stereo and motion are present. However, since color can be preattentively processed, it probably helps significantly in conditions like printed media where other depth cues are lacking.

3.3 Texture styles

Six different texture cases were used. They are shown in Fig. 4, with larger versions in Fig. 15. The (0) case is a single surface with a grid texture. This gives a base-line for minimum subject error. (1) is the no-texture case. (2) has equal spaced grids on top and bottom surfaces, (3) has a small top grid with a large bottom grid, and (4)

has a large top grid and a small bottom grid. All of the grid lines are drawn fully opaque, and in (2), (3) and (4) they are drawn at a width to provide 50% average opacity on the top surface. (5) was constructed from equal spaced grids on top and bottom, but with thinner lines on top and a translucent background between lines. The width of the thin lines was set to provided an average of 25% opacity, and the background transparency to give another 25%, making the total average opacity again 50%

Each subject was presented with a total of 240 probes; 40 total probes for each of the 6 texture types, with half of the probes on the top surface, and half on the bottom. The experiment was divided into 5 sets, with a random ordering of textures and top/bottom probes within each set to minimize learning and fatigue effects. The surfaces were generated randomly, and each probe was placed randomly on the surface.




Fig. 4. The six textures styles used for our first experiment.

3.4 Questionnaire

After the experiment, subjects were asked to order the 5 layered texture cases according to how ‘attractive’ they found each visualization. This was asked because a visualization that subjects find dreadfully ugly may not be useful even if by perceptual measures it works fine. Subjects were also asked to order their estimated ‘performance on the task’ for each visualization. Some subjects admitted they had trouble separating these two qualities, but most answered differently for ‘attractiveness’ and ‘performance’.

4 RESULTS

Probe error rates were analyzed using a two-way analysis of variance measuring the effects of texture type and subject. Figures 5-7 show multi-comparison plots with a controlled family-wise error rate of $\alpha=0.05$ showing each texture type and corresponding errors. Error is measured as degrees of angle between the probe direction and the correct normal direction. Line lengths show the 95% confidence intervals for each texture style, so lines that do not overlap can be

considered significantly different, according to Tukey’s honestly significant difference (hsd) test. Color variation and the dashed lines are included for clarity Red lines are all texture styles that are significantly different from the selected blue line. From Fig. 5 it is clear that there is a significant difference [$df=5$, $F=180$, $p=0$] in error between the single-surface case (0) the un-textured layered case (1) and the grid-textured layered cases (2-5). Looking at a single surface alone gives a relatively low error average of 9.33° . Although the textured layered surfaces have higher errors (12.4° , 11.7° , 11.2° and 11.1°), the best case was the thin lines (5) with an average error only 1.8° worse, or about 20% more error than for a single surface. It is also clear that the no-texture case (1) is by far the worst, with an average error of 22.1° . Simply guessing straight up for every probe would give an average error of 25.6° .

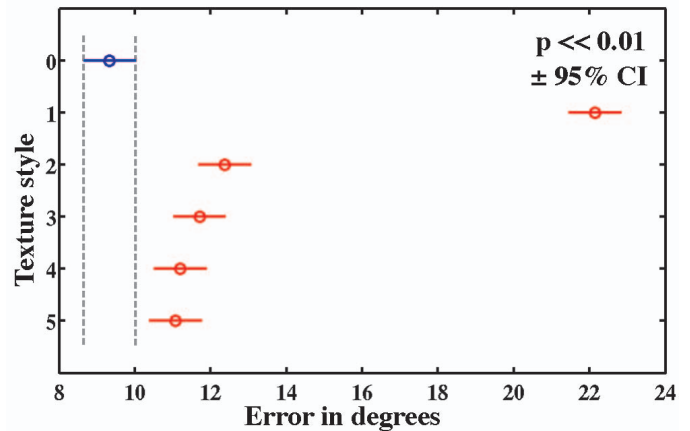


Fig. 5. Average error rates for all textures.

Next, we compared error rates for the four layered, grid-textured cases on top and bottom surfaces separately. Figure 6 shows no significant differences [$df=3$, $F=1.1$, $p=0.34$] between the grid styles for top surface error. However, Fig. 7 shows that (2) is significantly worse than both (4) and (5), while (5) is significantly better than both (1) and (2) [$df=3$, $F=5.5$, $p=0.001$]. The texture-subject interaction term was significant; several subjects performed better under the equal-grids or small-top grid conditions, but overall (4) and (5) are better. It seems that equally spaced grids with equal line widths make it most difficult to see the bottom surface.

We have two possible interpretations for these results. The first is a global argument, and the second a more local one. The global argument is that having grids with different spatial frequencies makes it easier for the human visual system to separate the visual signals across the two surfaces. The top and bottom frequency spectra in (5) are also different because the thin lines on the top surface add more high frequencies. Our second, local argument involves line junctions. As discussed earlier, humans use line junctions for reconstructing 3D shape from line drawings. But with equal grid spacing, as in (2), most of the bottom surface line junctions are at least partially occluded by the top grid. The same is true in (3), but the thicker bottom lines allow easier interpolation of the grid lines, so that junctions can be inferred. On the other hand, both (4) and (5) have larger openings in the top grid because of grid spacing and line width respectively. Consequently, at least one line junction is visible through every opening, improving the ability to extract 3D shape.

Interestingly enough, although subjects were told to try to keep a consistent pace while aligning the probes, time spent on probes mirrors the error rates. This effect is apparent in Fig. 8 [$df=5$, $F=5.9$, $p=0.00002$]. Although the effects were not as extreme, this does support the findings based on error-rates; it seems that subjects tended to spend more time and perform poorly on some of the visualizations. Unlike error, however, times for top and bottom surfaces were quite similar, and so are not shown separately.




Fig. 6. Top surface error rates for grid textures.

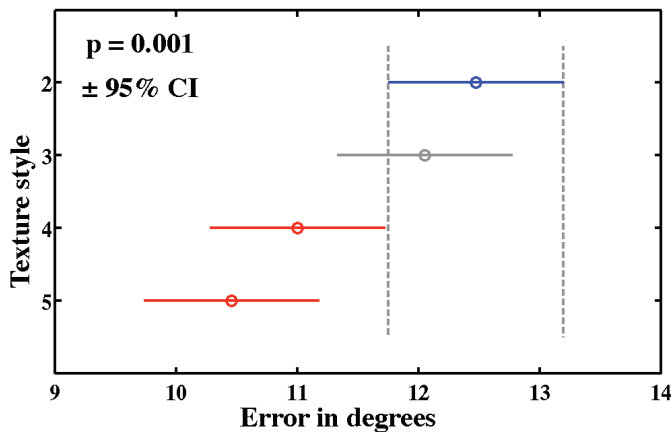


Fig. 7. Bottom surface error rates for grid textures.




Fig. 8. Time for all textures.

The results for the post-experiment questionnaire are shown in Fig. 9. Results for all 14 subjects were summed with a 'Best' answer being +2, 'Better' being +1, 'Average' being 0, 'Worse' being -1 and 'Worst' being -2. This makes +28 and -28 the best and worst possible scores. Clearly the subjects' beliefs about how well they could perform the tasks mirror their actual results. The top half of Fig. 9 shows performance ratings. Every subject rated the untextured case as 'Worst', and most thought that (4) and (5) were best. The subjective 'beauty' results shown in the bottom half of Fig. 9 strengthen the argument for using texture types of (4) or (5) as well. One goal of visualization is to be attractive; subjects are more likely to use a visualization technique if they find it attractive to look at.

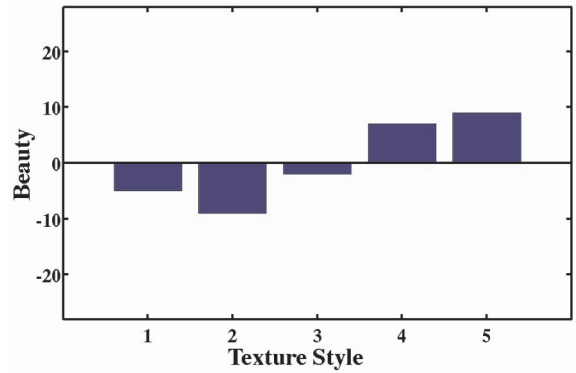



Fig. 9. Questionnaire results.

5 EXPERIMENT TWO

In our next experiment we investigated how variations in the opacity of the top surface affect error rates on the top and bottom surfaces. Opacity was varied between 30% and 70% on the top surface, with 5 levels. Grids were drawn with a combination of the (4) and (5) styles from the previous experiment. We hypothesized based on the results from the first experiment that thin lines and larger grid spacing on the top surface made it easier to see and interpolate line junctions on the bottom. Therefore the grids for experiment 2 were drawn with the larger spacing but with thin lines and a translucent background. Both line width and background opacity were varied to create the 30%, 40%, 50%, 60% and 70% opacities shown in Fig. 10. Close-ups of these are shown in Fig. 16.




Fig. 10. Top surface opacity levels of 30-70% from left to right.

Seven subjects ran this experiment (3 female, 4 male, 1 professor, 6 students). Three of these subjects had participated in the previous experiment. Conditions and training were like those of the first experiment. Each subject ran 200 total probes giving 20 probes each for both top and bottom surfaces for each of the 5 opacity levels. Results are shown in Figs. 11-13.

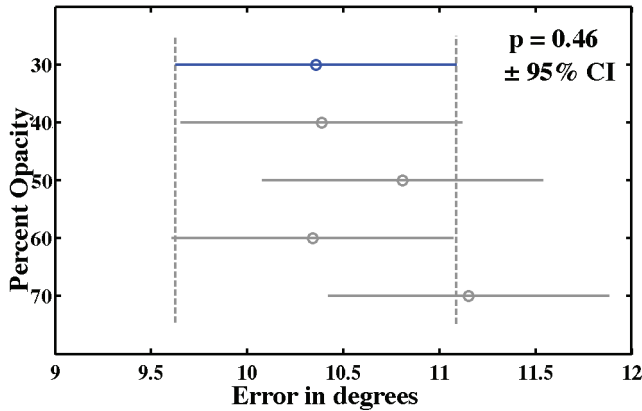


Fig. 11. Average errors for different top opacities.




Fig. 12. Bottom surface errors for different top opacities.

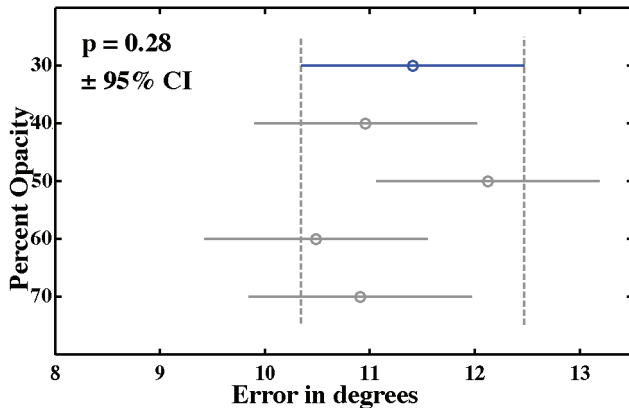


Fig. 13. Top surface errors for different top opacities.

Surprisingly, our results show almost no significant differences between the various opacities. Figure 11 shows average error for both top and bottom surfaces, with no significant effect [$df=4$, $F=0.9$, $p=0.46$]. Figure 12 shows a general trend toward higher bottom error with higher top surface opacity [$df=4$, $F=2.7$, $p=0.03$]. This we would expect. On the other hand, accuracy on the top surface, in Fig. 13, did not vary significantly with opacity [$df=4$, $F=1.3$, $p=0.28$]. This is surprising, but it fits with our results for experiment 1, in which top surface errors had no significant variation with texture spacing. We think that perhaps seeing the top surface is generally the easier task, as it does not require estimating surface shape from behind occluding objects. Also, the grid seems to work very well for showing the top surface shape even with very little opacity.

6 FOLLOW UP EXPERIMENT

We were curious to see if our hypothesis that a combination of larger top grid and thinner top lines would work better than changing grid spacing and line width individually. A small experiment was run identical in setup to the first two experiments, comparing (4) and (5) from the first experiment with the 50% opacity case from the second experiment. 4 subjects ran the experiment (3 male, 1 female, 1 professor, 3 students), though 3 had participated in both previous experiments, so only one subject was truly naïve to the purpose of the study. Possibly due to the small number of subjects, no significant differences were found in either error or time differences between the three textures styles on either top or bottom surfaces. However, as suspected, the trend on the bottom surface was for the thin lines to perform better than the larger grid spacing, and the combination of the two performed the best. Also, although still not significant, the trend on the top surface was the opposite. The combined thin lines, large spacing grid performed worst on average, while the large spacing grid performed the best. Again, time spent correlated strongly with error rate. Although no significant results were found, this seems to demonstrate the balancing act required to show both surfaces clearly, with equally spaced grids and thin top lines being (arguably) the best solution.

7 DISCUSSION

The most simple, clear result from our experiments is that well-chosen textures give far better results for surface shape determination than shaded surfaces. Other researchers have shown this result, but we emphasize it here because, in spite of this, simple shading is still the most common method used in layered surface visualization. When viewing our simple shaded case, several subjects said they had difficulty telling whether the probe was on the top or bottom surface, or even which surface was on top or bottom. Many also said that the probe appeared to be floating in midair, not 'attached' to any surface. Without fine-grained texture, the combination of stereo, motion and shading were apparently not enough to allow subjects much better than a random guess. To be fair, common layered surface visualizations often use objects with boundaries and/or differing spatial-frequencies, allowing shading to act like a texture. This lessens the confusion and decreases the likelihood that subjects will perceive a feature on one surface that is actually on the other surface. Still, we found compelling evidence for including more use of texture in layered surface visualization.

Also, our results support an argument for using layered surfaces in the first place. Often, visualizations that could include layered surface techniques simply display the surfaces such that the top surface occludes nearly all of the lower ones, or at a steep angle arranged vertically so that little actual overlap occurs. Although this certainly conveys the surface ordering and some information on surface shapes, occlusion of lower surfaces hides information that might be useful, and steep angles compress and hide information within each surface. Results from our best texture case have only 20% more error while showing twice the information of the single-surface case. Certainly, for some applications a decrease in accuracy would be a fair trade-off given the extra information shown. Even visualizations with more than two layers might benefit from showing two of those surfaces simultaneously, letting the others be occluded or made mostly transparent.

Our results also confirm the hypothesis that for the task of estimating surface shape, larger top surface texture patterns or thinner top lines enhance perception of the bottom surface. That is, grids with larger transparent areas make it easier to see the bottom surface without compromising top surface visibility. We were surprised to find how little top surface opacity affected surface shape estimation accuracy over a large range of different opacities. Like the texture spacing, opacity only seemed to affect bottom surface accuracy, with low and middle opacities all working quite well. We feel strongly that the small effect opacity appears to have is due to the initial quality of the visualization, including lighting, viewing

angle, texture, stereo and motion. We believe that by setting these parameters carefully, a poor choice of opacity simply has less of an effect. Also, it should be noted that for high top opacities there was a significant drop in accuracy, so for applications where the bottom surface is more important, such as the brain inside the skull, care should be taken in the opacities used. On the other hand, it is intriguing to see that the top surface accuracy can be very high even with very low opacities. A low opacity top surface distracts little from the bottom surface, but apparently conveys a great deal about its shape. However, we would hesitate to recommend using extremely low top opacities for tasks that require picking out features quickly, since in this case shading may be a more important factor than texture, and low opacities reduce shading information.

8 VOLUME VISUALIZATION EXAMPLE

In our experiments, we analyzed how various texture styles affect visualization of layered surfaces. However, a great deal of visualization deals with volumetric data rendered using volumetric techniques rather than as isosurfaces. Here we develop an example of how textures might be incorporated into segmented volume visualization.


Figure 14 shows various renderings of a model we created of a horse's hoof, which consists of an outer hoof wall, and an inner coffin bone. The polygonal models were created in Maya, and rendered using the Renderman Shading Language. The original purpose of the model and visualization was for instruction: to show the shapes and relationship between the hoof wall and coffin bone.

First, in Fig. 14a we show surface rendering of the hoof. The polygon mesh is shaded using Phong shading and transparency. Only forward-facing polygons are rendered. The scene is lit with two spotlights and an ambient light, and colors are chosen to roughly correspond to actual hoof wall and bone tissue. However, to someone unfamiliar with the hoof structure, the visualization is confusing because the opacity does not generally correspond with the actual thickness of the surface.


The volume-shading example in Fig. 14b improves visualization of the isosurfaces by rendering opacity proportional to thickness along the viewing vector. Also, edges were enhanced by varying color from lighter to darker depending on the dot product between the surface normal and the viewing vector. The structure and thickness become much clearer in this example. Finally, Fig. 14c shows an example using grid texturing. Here, opacity is determined by the union between thickness and the grid texture. This tends to make the texture more prominent in thin parts of the volume. We feel that this is the most successful visualization at showing shape clearly. Although our results are for texturing of surfaces under conditions of motion and stereo, we feel that careful texturing according to our results can improve visual clarity even for the case of printed media. However, it should be noted that for these types of surfaces the choice of texture mapping can dramatically effect results. Since the surfaces are not conducive to planar texture projections, they require either hand mapping, or the methods such as those by Gorla *et al.* [10] for texture mapping of arbitrary surfaces.

9 FUTURE WORK


Certainly, these experiments only scratch the surface in terms of possible variations in texture patterns on layered surfaces. The choice of a grid, while not arbitrary, could certainly be varied, and our group is working on studies to determine if results about random, line and grid textures on single surfaces also apply to layered surfaces. Also, applying the work we present on texture style to the work on principal-direction textures for layered surfaces might prove quite fruitful. Finally considering our results on low-opacity top surfaces, we feel that investigating techniques for visualizing three or more surfaces might be surprisingly successful.



a) surface shading



b) volume shading



c) textured volume shading

Fig. 14. Horse hoof shading examples.


ACKNOWLEDGMENTS

The experimental design, including the probe, surface structure, and stereoscope was by Colin Ware. We would like to thank the Texas A&M Vizlab faculty and students for participating in our


experiments, Dr. David Hood for his help in understanding the structure and shape of a horse hoof, and Mayank Singh for his assistance and suggestions. We would also like to thank the anonymous Visualization 2007 reviewers for making several recommendations, which resulted in significant improvements to the paper. This work was supported in part by the National Science Foundation ITR's 0326194 and 0324899, the Center for Coastal and Ocean Mapping – University of New Hampshire, and the Visualization Laboratory – Texas A&M University.

REFERENCES


- [1] D. Acevedo and D. H. Laidlaw. Subjective Quantification of Perceptual Interactions among Some 2D Scientific Visualization Methods. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization)*, 12(5), Sept-Oct 2006.
- [2] A. Bair, D. House, C. Ware. Perceptually optimizing textures for layered surfaces. *Proceedings of Symposium on Applied Perception in Graphics and Visualization*, pp. 67-74, 2005.
- [3] A. Bair, D. House, C. Ware. Texturing of Layered Surfaces for Optimal Viewing. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2006)*, 12(5), 1 pp.125-1132, 2006.
- [4] H. Barrow and J. Tenenbaum. Interpreting Line Drawings as Three-Dimensional Surfaces. *Artificial Intelligence*, 17, pp. 75-116, 1981.
- [5] R. Bhatt and E. Bertin. Pictorial Cues and Three-Dimensional Information Processing in Early Infancy. *Journal of Experimental Child Psychology*, 80 pp. 315-332, 2001.
- [6] M.J. Black, R. Rosenholtz. Robust Estimation of Multiple Surface Shapes from Occluded Textures *International Symposium on Computer Vision*, pp. 485-490, 1995.
- [7] B. Cummin, E. Johnston, A. Parker. Effects of different texture cues on curved surfaces viewed stereoscopically. *Vision Res.* 33(5-6):827-38, 1998.
- [8] J. Enns and R. Rensink. Preattentive Recovery of Three-Dimensional Orientation from Line Drawings. *Psychological Review*, 98(3), pp. 335-351, 1991.
- [9] H. Feichtinger, T. Strohmmer. *Gabor Analysis and Algorithms: theory and applications*. Birkhäuser 1998.
- [10] G. Gorla, V. Interrante, G. Sapiro. Texture Synthesis for 3D Shape Representation, *IEEE Transactions on Visualization and Computer Graphics*, 9(4), pp. 512-524, 2003.
- [11] D. House, and C. Ware. A method for the perceptual optimization of complex visualizations, *Proceedings of Advanced Visual Interfaces (AVI '02)*, 148-155, 2002.
- [12] D. House, A. Bair, and C. Ware. On the optimization of visualizations of complex phenomena, *Proceedings of IEEE Visualization 2005*, pp. 87-94, 2005.
- [13] D. House, A. Bair, and C. Ware. An Approach to the Perceptual Optimization of Complex Visualizations, *IEEE Transactions on Visualization and Computer Graphics*, 12(4) pp. 509-521, 2006.
- [14] 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, pp. 197-206, 2002.
- [15] V. Interrante, S. Kim. Investigating the Effect of Texture Orientation on the Perception of 3D Shape. *Human Vision and Electronic Imaging VI*, 2001.
- [16] V. Interrante, H. Fuchs, and S.M. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Trans. on Visualization and Computer Graphics*, 3(2) pp. 98-117, 1997.
- [17] V. Interrante and C. Grosch "Visualizing 3D Flow", *IEEE Computer Graphics and Applications*, 18(4), pp. 49-53, 1998.
- [18] S. Kim, H. Hagh-Shenas, and V. Interrante. Conveying shape with texture: experimental investigations of the texture's effects on shape categorization judgments. *IEEE Trans. on Visualization and Computer Graphics*, 10(4) 471-483, 2004.
- [19] 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.
- [20] D. Knill. Perception of surface contours and surface shape: from computation to psychophysics. *J. Opt. Soc. Am. A*, 9(9), pp. 1449-1464, 1992.
- [21] J. Koenderink and A. van Dorn. Affine Structure from Motion. *Journal of the Optical Society of America A*, 8(2) pp. 377-385 1991.
- [22] D. H. Laidlaw, M. Kirby, S. Davidson, T. Miller, M. DaSilva, W. Warren, and M. Tarr. Quantitative Comparative Evaluation of 2D Vector Field Visualization Methods. In *Proceedings of IEEE Visualization 2001*, pp. 143-150. IEEE, October 2001.
- [23] P. Mamassian, D. Kersten. Illumination, shading and perception of local orientation. *Vision Research*, 36(15), pp. 2351-2367, 1996.
- [24] V. Ramachandran. Perceiving Shape from Shading. *Scientific American*, 259 (2), Aug. 1988.
- [25] J. Saunders, B. Backus. Perception of surface slant from oriented textures. *Journal of Vision*, 6, pp. 882-897, 2006.
- [26] G. Sweet, C. Ware. View Direction, Surface Orientation and Texture Orientation for Perception of Surface Shape. *Proceedings of the 2004 Conference on Graphics interface*. ACM International Conference Proceeding Series, vol. 62, pp. 97-106, 2004.
- [27] J. T. Todd, and E. Mingolla. The perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 583 595, 1983.
- [28] J. T. Todd and F. Reichel. Visual Perception of Smoothly Curved Surfaces from Double-Projected Contour Patterns. *Journal of Experimental Psychology: Human Perception and Performance*. 16(3) pp. 665-674, 1990.
- [29] J. Todd, A. Oomes, J. Koenderink, A. Kappers. The Perception of Doubly Curved Surfaces from Anisotropic Textures. *Psychological Science*, 15(1), pp. 40-46, 2004.
- [30] T. Urness, V. Interrante. E. Longmire, I. Marusic, S. O'Neill, T. Jones. Strategies for the Visualization of Multiple Co-located Vector Fields. University of Minnesota Tech. Report 05-032. 2005.
- [31] C. Ware, and G. Frank. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics* 15(2): pp. 121-140, 1996.
- [32] C. Wheatstone. Contributions to the physiology of vision. Part the first. On some remarkable and hitherto unobserved phenomena of binocular vision, *Philosophical Transactions of the Royal Society*, 128, pp. 371-394, 1838.




(0) Single surface with grid texture




(1) Layered surface, no texture




(2) Equally spaced top and bottom grids



(3) Small top grid, large bottom grid



(4) Large top grid, small bottom grid



(5) Equally spaced grids with thin top lines, and translucent top.

Fig. 15. Full-screen views of the 6 textures styles used for the first experiment. Probes are shown on the bottom surface.




Fig. 16. Close-up of opacity levels for the second experiment.