9. CONCLUSIONS

9.1 Summary and Conclusions

The purpose of this line of research is to investigate layered surface texturing. Visualization of layered surfaces can be useful both for scientific exploration of layered data, and for instructional purposes like communicating well-understood layered data to other scientists and the general public. Showing layered data simultaneously should allow viewers to make connections between the data layers that would be more difficult to see if the data were presented separately. Texture was known to be a strong cue for surface shape, and has been shown in these studies to be an important cue for showing layered surfaces as well. The studies presented here used a broad to narrow approach to find what factors are important for layered textures, what values are optimal, and what interactions are important between the factors. Each study used experiments run on human subjects to find optimal texture conditions for layered surface visualization. Each experiment gave subjects a task to complete for a set of textures that varied according to a given texture parameterization. The initial experiments were broad, using many parameters, while the later experiments were more specific, varying only a few select texture parameters.

There were two main types of task used to measure perception of the layered surfaces. The task used in most of the experiments was a measure of surface direction perception; subjects had to align a probe to the surface normal at a particular location. However, one experiment used a feature-finding task; subjects had to identify bumps on a noisy surface. Optimal textures were found to vary depending on the task.

Texture was found to be extremely important for the surface normal estimation task. In a study that compared textured layered surfaces to un-textured but shaded layered surfaces, the errors on the un-textured surfaces were almost as bad as a random guess, far worse than any of the texture conditions studied. Overall, it was found that clear textures with strong directionality, such as grids or lines are the most important for showing shape information.

On the other hand, texture seemed less critical for the feature-finding task. In this case, optimal textures were subtle, with very high luminance. In most cases, the best textures had only one surface with a noticeable texture, leading us to suspect that texture was primarily being used to distinguish surfaces. Also, a bright texture allows for clear shading information: a strong cue for identifying features on a surface. Therefore, if a visualization is being used simply for identification of important features, textures should be subtle and bright.

Since surface normal estimation is a more difficult task than feature finding, and probably gives a better estimate of shape understanding, it was used for the finer-grained studies of texture optimization. Texture opacity was found to have an unexpectedly wide range of usable values; in a two-layer visualization, the top surface opacity could vary from 30% - 60% without any significant error increase. Texture sizes had a more complicated effect. The pilot studies suggested that larger top textures seemed to work better. Two studies explicitly varied texture size. Results from these showed that top texture size had a significant effect, but it involved a trade-off between performance on the top and bottom surfaces. Larger openings in top texture grids make the bottom surface easier to see, but the top surface harder to see.

A feasibility experiment was run to estimate how many surfaces could be layered before the visualization is too complicated for subjects to get any worthwhile information. Results from this study suggest a limit of 5-6 surface layers, at which point the visualization is too complicated to understand. Probably 3-4 layers are the most that could be practically shown at once. User interaction could possibly increase this limit. See the discussion for more ideas on user interaction for multiple layers.

Textures with lines in the principal curvature direction of a surface have been thought to be optimal for surface shape perception, but most of the research has focused on developable surfaces, a very limited set of all possible surface shapes. Therefore, our final experiment compared projected grid textures with principal direction textures. On average, projected grid and principal direction textures were found to work equally well for shape perception. However, the projected grid textures were worse in areas where the surface normal pointed toward the viewer and significantly away from the projection direction. Overall, subjects used texture line information as expected for estimating the surface direction. Errors were lowest for edge-on surface patches, and highest for front-facing surface patches. A rotational bias was found on a per-subject basis. If the probe rotation can be thought of in terms of latitude and longitude, where the north pole corresponds to the up viewing vector, most subjects tended to over-rotate the surface normal probe toward the equator. The magnitude of this over-rotation was in proportion to how close the probe was to the equator. This suggests that steps could be taken to reduce the overall error by minimizing this bias.

9.2 Discussion

Only rather simple textures were used in this line of research. These included projected grids, dots and hatches, and principal direction textures made with discrete marks. Many other texture types are available, and might prove useful for shape perception. Among these are surface-dependent textures and viewpoint-dependent textures. Principal-direction textures are an example of surface-dependent textures, but others might include textures based on the surface height, curvature magnitude, or the direction of the surface relative to another point such as a light source. One example of a viewpoint-dependent texture is an edge enhancement algorithm. Surface boundaries are a strong shape cue. That is, surface shape on a boundary is more constrained than in other areas, and the shape seen at the boundaries can be interpolated to improve shape perception elsewhere. Silhouette shaders are already commonly used in non-photorealistic rendering. To make these shaders work for an interactive surface visualization, they would need to render fast enough for interactive exploration of the scene. Also, the algorithm would need to have temporal coherence. Since edge and boundary lines are viewpoint specific, the lines that are drawn would change as a user rotates a camera to explore a surface. A good algorithm would ensure that no unnatural temporal jumps in the lines would occur to distract the user. Since humans are very sensitive to motion, any abrupt changes in the overall line shape and location would probably make the method unusable. But if the temporal coherence issues could be overcome, view direction-dependent textures could probably provide stronger surface shape cues. Ideally, an optimal texture would be one that finds optimal surface and viewpoint-dependent locations for showing shape information, adds clear and bias-free texture marks in those locations, and leaves the rest of the surface empty to allow space for information from other surfaces room to show.

The methods used to measure shape perception in this line of research were relatively simple. We included subjective ratings of surfaces, feature finding tasks and surface normal estimation. Surface normal estimations were able to measure not only magnitude, but also direction of error, which led to several interesting findings. However, one thing that has not been measured directly is how the subjects are actually using texture information. Eye-tracking is one method that could show what information subjects are actually focusing on when looking at the surfaces. For example, it would be interesting to learn how much surface area around any probe was actually looked at to perceive the surface. This might vary depending on the task. Certainly feature finding would probably involve a quick scan of the entire surface, with longer focus time spent on each feature. We hypothesize that surface normal estimation would only use the several closest line junctions to estimate a direction. Curvature estimation might require more data: synthesis of the global properties of the surface, or it might simply use the local curvature of a few line junctions.

Curvature has been measured before using a forced-choice categorization task, where subjects had to say if a surface was concave, convex, cylindrical or saddle-shaped. Subjects have also been asked to draw curves to match the perceived curvatures of singly curved surfaces. However, it would be interesting to find a way to get a quantitative measure of perceived local curvature for arbitrary surfaces. One possible way to do this might be asking the subject to align several probes in an area, and taking the difference in surface normals as a measure of the perceived curvature. However, since we don’t yet have a good understanding of how much local vs. global information is used in estimating surface normals, it would be hard to know if the probe directions were being set relative to each other based on curvature perception, or independently based on surface normal perception at each location. Still, since curvature can be derived from change in direction, and vice versa, such an experiment should give useful results.

A second possibility would be to develop a ‘curvature probe’ similar to the surface normal probe to directly measure perceived curvature. For example, at a given surface location, the subject could manipulate a small surface patch to look like the actual surface. The subject could vary two curvature magnitudes and a rotation to try to match the surface patch with the real surface. This would measure the perceived minimum and maximum curvature directions, as well as their magnitudes. Since the surface patch direction would be set to correctly match the actual surface normal direction, biases in surface direction would not affect the curvature estimations. We note based on experience, that if stereo cues are used the surface patch should be translated away from the actual surface. This is because subjects can be very accurate at matching the relative heights of the surface patch with the real surface even without accurate perception of the surface curvature or direction. Thus care should be taken to make sure that the surface patch is seen from the same camera angle and field of view, despite being positioned away from the actual surface. Also, the surface patch should have visible edges to maximize the understanding of its curvature.

Surface normal and curvature perception both measure the perception of a single surface, no matter how many surface layers are being displayed at a particular time. However, the reason that multiple surfaces would be shown in a single scientific visualization is to find relationships between the surfaces. Relationships might include whether there are correlations between the directions or curvatures of the surfaces, or identifying locations where the surfaces are very close together or far apart. Finding relationships might work well as a feature-finding task, where subjects are given a pattern or relationship to look for in the surfaces. However, under the assumption that subjects are able to perceive the individual surfaces accurately, finding relationships becomes a problem of understanding the relative locations of points on each surface. Therefore, probes to measure perceived direction and distance between random points on each surface might be a good quantitative measure of relationships between surfaces.

Monocular, static images of layered surfaces are much more difficult to see than visualizations with stereo and motion cues. Although we have done no studies directly comparing visualizations with and without the stereo and motion cues, subjects averaged less than 10° of error for a single surface in a study that used both cues, while subjects had an average error near 13° in a similar study that used only stereo cues. Since static images have fewer cues to surface shape, we would expect that static, monocular viewing conditions would have larger errors overall, and any biases would probably be exaggerated. Also, some parameters that were not particularly important for moving, stereo images might be more important without those cues. For instance, choosing an optimal viewpoint is probably quite important, because viewers do not get a variety of viewpoints. Also, differentiation of different texture layers using texture parameters like hue are also likely to be more important. Stereo and motion cues are very helpful for separating one surface layer from another. Therefore, cues that distinguish one surface from another should be useful for static images.

Not all three-dimensional data naturally fits a surface model. For example, in medical data, because of the nature of the scanning technology, most data is in the form of volume data. However, most of the tissues and organs of the body do have well-defined surfaces. Careful segmentation algorithms can successfully extract these surfaces from the volume data for texturing. There is no reason why volume rendering and surface rendering methods cannot be combined. Volume rendering can change depending on the transfer function. The transfer function is used to transfer the volume data along a ray direction into a color at a pixel. One of the most common transfer functions simply adds the densities at points along a ray, so that the image shading is an indication of the object thickness and material type. Treating the colors produced from volume rendering as surface shading means that they can be combined with texturing on the object’s surface. Therefore, there is no reason why texturing cannot be used equally well with surface and volume data.

The subjects in our experiments had relatively little ability to interact with the surfaces they were seeing. Although the surfaces rocked in most of the experiments, the subjects could not move around as they wished, and the camera was at a fixed distance from the surfaces. Many types of interaction might be added to improve overall perception of the surfaces. Certainly, allowing the users free movement within the scene would improve surface perception because the users could move to optimal camera angles depending on the shape and direction of the surface. Turning a difficult-to-see shape so that it shows as a silhouette would allow subjects to better interpret the 3D shape in that area.

Also, selective display of surfaces and parts of surfaces could help subjects understand single surfaces or pairs of surfaces at a time, so that a more complete mental picture of all the surfaces can be built. Based on the difficulty of seeing six surfaces together in the feasibility study, layering of a large number of surfaces is not likely to work without significant user interaction. A simple interface to turn the display of individual surfaces on and off would let users inspect any combination of layers at a time. This would require a user to remember the surface shapes and relationships between selected surfaces. However, we guess that the most important relationships are often between individual pairs of surfaces, so this simple method might work.

A slightly more complicated method could allow users to vary the opacity of portions of surfaces. That is, users could select areas of interest where they want to see lower level surfaces clearly. This portion of the upper level surfaces would become fully or mostly transparent so the user could see through them, but the rest of the upper surfaces would still provide context. This is similar to the artistic methods of using selective opacity only where the texture significantly helps show the upper surface, or where the lower surface is not important, and could be either user-driven, or possibly automatic.

Interpenetrating surfaces are another problem that was not considered in this line of research. Surfaces that have no clear ordering present several interesting problems for both the implementation and the theory of visualization. For one, if the surfaces are represented as polygonal meshes, a layered visualization must be drawn from the most distant polygons at any pixel to the closest polygons for the composition operations to work correctly. This means that intersections of the surfaces must either be pre-computed, or the polygons must be sorted on a per-fragment basis. For surfaces based on distance fields or other functions, this is less of an issue.

However, the implications for how to texture interpenetrating surfaces are more interesting. Certainly, the textures should be consistent enough over each surface so that the continuity of each surface is clear. Yet, the bottom-most surface will likely be fully opaque, while any upper surfaces will be partially transparent. It would be interesting to test whether an abrupt or a smooth transition between the opacities works better. Based on the importance of silhouette and contour lines for perception, we anticipate that emphasizing the interpenetration contour might be useful for shape perception. Other style parameters might vary between the surfaces, such as color, size, frequency, line direction and texture type. Since several top-layer factors have been shown to affect bottom-surface perception, it would be important to consider which texture parameters were consistent along a particular surface, and which were consistent depending on the layer depth.

All of this research has been on showing surface shape, but other types of data might exist on the surfaces, and using texture to show that data could allow users to see relationships between the data and the surfaces themselves. Mapping data onto the textured surfaces is an interesting problem. Since surface shape is shown best with directional textures, any data mapping would have to work with the directional texture. Data is often shown with glyphs, using variables like shape, color, density and direction of glyphs to encode data visually. One workable method might be to have texture shape marks (probably lines of some sort), and glyph marks on the same surface. However, changes in glyph compression, size and density could be interpreted as changes in surface angle, or distance from the camera. Therefore care would have to be taken that the effects of viewpoint and surface direction did not interfere with interpretation of the data encoding. Hue would actually be a good glyph parameter. A single dimension of data on each surface could be encoded using different hues for each surface. Mark blur, and line parallelism are other parameters that are independent of viewpoint and surface direction. Symbols would certainly also be useful for marking feature locations.

Color, blur and line parallelism are a rather restrictive set of glyph parameters. We believe that if the surface shape texture has strong regular distance cues, such as would be found in a regular projected grid, then glyph size, shape, density and direction could also be used to encode data. With the strong shape cues provided by grid contours, subjects are likely to correctly interpret shape changes as changes in the actual glyph shape rather than a rotation of the surface itself. Motion and stereo cues would help support this interpretation so that subjects could correctly read both the surface shape and changes in data glyphs.

Given how important texture has been found for surface shape perception, it is surprising that it is hardly ever used in scientific applications. Certainly care must be taken that textured do not obscure or confuse the surface data being shown, and scientists are not always known for their artistic abilities. However, we feel that with the guidelines presented in these studies, even people unfamiliar with visualization principals could successfully use texture to show surface shape. One use for texture is medical visualization. In many cases, disease diagnosis is dependent on feature finding of abnormalities, such as tumors. For these tasks, any textures used should be very subtle so that the texture does not interfere with any feature data. However, in some cases the shape is very important, such as determining the size and shape of a tumor to be surgically removed. In this instance, we would strongly recommend the use of texture to improve shape perception. Like any tool, texturing must be used with care, but when used to best advantage, texturing can provide huge benefits to the complicated process of shape perception.