Opacity-modulating Triangular Textures for Irregular Surfaces

Penny Rheingans
Department of Computer Science
University of Mississippi

Abstract

Many scientific and medical visualization techniques produce irregular surfaces whose shape and structure need to be understood. Examples include tissue and tumor boundaries in medical imaging, molecular surfaces and force thresholds in chemical and pharmaceutical applications, and isosurfaces in a wide range of 3D domains. The 3D shape of such surfaces can be particularly difficult to interpret because of the unfamiliar, irregular shapes, the potential concavities and bulges, and the lack of parallel lines and right angles to provide perspective depth cues. Attempts to display multiple irregular surfaces by making some or all of them transparent further complicates the problem. Texture can provide valuable cues to aid in the interpretation of irregular surfaces. Opacity-modulating textures offer a mechanism for the display of multiple surfaces without the extreme loss of clarity of multiple transparent surfaces. This paper presents a method for creating simple repeating textures and mapping them onto irregular surfaces.

1. Introduction

Many applications of visualization require the display of irregularly shaped surfaces. Understanding terrain undulations, cloud shapes, molecular surfaces, or tissue masses requires the interpretation of surfaces with irregular, and often unfamiliar, shapes. Display of the 3D statistical surface defined by scalar values in a 2D domain, commonly called a mountain plot, can involve similar irregular surfaces. Additionally, visualization operations often create additional irregular surfaces for interpretation. Isosurface generation is the most common of these, defining the boundaries of 3D regions in a wide variety of applications. Isosurfaces are used to delineate 3D regions where certain conditions hold, such as places where air pollutant levels exceed legislated limits, tissue density deviates from normal values, toxin levels in river sediment are unacceptably high, or air flow creates high pressure values. In such situations, observers use these isosurfaces to judge whether additional pollution restrictions should be imposed, a biopsy should be performed to investigate a suspected tumor, the river bottom should be dredged to remove contaminated sediment, or the building should be redesigned to improve air flow. Such critical decision-making depends on accurate interpretation of isosurface shapes.

Author's address: Department of Computer Science, University of Mississippi, University, MS 38677. Email: rheingans@cs.olemiss.edu.

Arbitrary surfaces are particularly hard to understand because they can lack many of the visual cues observers generally use to interpret three-dimensional shape. The lack of parallel lines and right angles reduce the power of perspective depth cues. Unfamiliar shapes limit interpretation based on past experiences with similar shapes. Irregular surfaces can self-shadow and self-obscure themselves in complex ways, reducing the effectiveness of lighting and obscuration cues. When surfaces are made transparent, in order to show what lies behind or within them, their shape becomes even more difficult to perceive accurately.

The application of texture to irregular surfaces can make them easier to interpret by providing additional shape cues. Texture can help disambiguate obscuration cues by clarifying how different parts of the surface meet and hide one another. The effects of perspective, and from them additional depth cues, are visible in the texture gradient. Specifically, the apparent scale of the projected texture decreases as the surface recedes from the observer. The texture in distant patches appears to be compressed, a denser pattern with smaller features. These perspective effects are likely to be more obvious on the texture than on the object itself. This is because texture shape in one area can be compared to texture shape in another, with differences indicating depth differences. Similar comparisons among geometric features of the object would not be as enlightening because object shape in different areas cannot necessarily be expected to be similar. The texture gradient also provides information about the orientation of surface segments, as textures undergo convergence on surfaces that drop away from the observer. More information about surface perception from texture can be found in the perception literature [for instance, Gibson50, Cutting84, Todd87, Cumming93].

Adding texture to transparent surfaces results in even greater perceptual benefits, since perception of shape from completely transparent surfaces is so difficult to begin with. For maximum benefit, the texture should combine transparent parts which show what lies behind with opaque parts which improve shape perception. I call such textures *opacity-modulating*, since they specify variations in the opacity of the surface as well as, or instead of, the color. The visual features of such textures can be either direct changes in color and intensity or indirect changes that occur when the identity of the visible surface changes as a result of changes in opacity.

2. Related Work

Schweitzer first proposed that artificial textures, those not necessarily mimicking any realistic texture, could be used to improve the perception of surface shape from rendered images [Schweitzer83]. He generated screen space texture elements with feature density and shape determined by the surface depth and orientation.

Previous uses of transparent textures to improve shape perception have included screen-space textures for automatic illustration [Dooley90], solid grid textures in medical imaging [Levoy90], and solid equation textures for mathematical visualization [Wejchert92]. Interrante [Interrante95] addressed the problem of accurately communicating the shape of transparent skin surfaces by adding opaque ridge and valley lines reminiscent of artists' renderings. This technique works very well on familiar shapes with relatively sparse ridge and valley features, such as the medical applications for which it was developed. The unfamiliar shapes and multiple self-obscurations of arbitrary surfaces limit the usefulness of ridge and valley lines.

Reaction diffusion textures, such as those developed by Turk [Turk91], can be generated directly on a surface, guaranteeing that the texture will follow surface contours. Unfortunately, systems of reaction-diffusion equations are generally tricky to define and slow to compute. Although texture values can be stored directly as vertex colors once they have been computed, making texture mapping solely a preprocessing step which does not impact view-time rendering speed, this forces texture repeats to be several times the size of typical polygons if texture details are to be visible.

In molecular visualization, texture mapping has increasingly been used to accelerate the mapping from molecular properties to surface characteristics [Teschner94]. In this property-based texture mapping, molecular properties at vertices are used to generate texture coordinates, so surface texture follows the distribution of property values across the surface. Duncan and Olson [Duncan95] describe location-based texture mapping for parametric molecular surfaces, characterized by a one-to-one correspondence between surface points and texture elements. In both sorts of molecular texture mapping, the objective is the display of additional information, rather than the improvement of surface shape perception. The textures used are generally opaque.

3. Requirements

What characteristics should a texture have in order to give strong shape cues on irregular surfaces? Appropriate textures are fine-grain repeating patterns which follow the surface of the object. Small, regular textures provide visual features without drawing undue attention to themselves. They must be able to be tiled without visible seams, since such seams are artifacts of the polygonal decomposition or specific texture mapping and might draw the observer's attention without conveying any useful information about surface shape. These stringent pattern requirements practically preclude that suitable textures can simply be scanned in; they need to be generated algorithmically.

Additionally, textures should be easy to generate and quick to render. Generally, this means that textures should be *image-based*, a precomputed image stored in texture memory, rather than *procedural*, defined by a procedure evaluated during the shading calculations made at each pixel. This allows applications to exploit the texture-mapping hardware available on common graphics workstations. Rendering speed is particularly important, since interactive viewing of complex three-dimensional shapes is a source of strong shape cues. *Kinetic depth effect* occurs when an object is rotated or translated to reveal previously obscured portions. The manner in which the surface appears from, and disappears to, behind the surface reveals much about the 3D structure underlying the 2D projection. Conversely, *head-motion parallax* results when the observer's virtual position in the environment changes, revealing different parts of the object. This position can be controlled in a variety of ways, but the strongest effects occur when head position is tracked to automatically control view position [vanDamme94].

Mapping a two-dimensional surface texture onto an arbitrary surface in 3D poses a problem. Arbitrary surfaces are not *developable*; that is, they cannot be laid out flat without distortion. By way of example, a cylinder is developable; a sphere is not. The difficulty lies in assigning texture coordinates to vertices in a way that accurately represents the surface geometry. A common method of mapping a texture onto a surface is to project the texture orthogonally or cylindrically onto the surface. The more the object is like a plane, or a cylinder, the better this works. Unfortunately, in these sorts of projections depth information (or the distance from the surface to the texture plane) is ignored. This distorts the texture in areas where the surface is not orthogonal to the direction of projection. Since we wish to use the apparent texture gradient distortions from the viewing and perspective transformations to convey shape information, texture distortions from the mapping process are clearly undesirable.

Alternatively, solid texture could be used to circumvent the problems of 2D-to-3D texture mapping. The appropriate texture value at each location on the object surface would simply be the value at the corresponding location in the 3D texture map. The result can be thought of as an object of the specified shape carved out of a block of the material described by the 3D texture. Unfortunately, solid textures do not provide the desired shape cues since the texture does not follow the surface contours. While texture gradient changes due to perspective will still be apparent, those due to orientation will not. In order to provide the appropriate shape cues, we require a texture that is generated directly on the surface; that is, one which is based on the shape of the surface itself, not on the underlying coordinate system.

4. Triangular Texture Space

A certain class of textures generated in a triangular texture space can greatly simplify the problem of mapping a 2D texture onto an arbitrary polygonal surface. Assume the surface is composed entirely of triangles. If not, an arbitrary polygon mesh can easily be transformed into a triangle mesh. For simplicity, use extreme points of (0,0), (1,0) and (0.5, $\sqrt{3}/2$) to create an equilateral triangular texture element. Now the problem has become that of mapping one triangle onto another. The assignment of texture coordinates to triangle vertices has been simplified to the assignment of the coordinates of one of the texture element extreme points to each triangle vertex.

Next, if the sequence of texture values along each side is identical, and consequently the values at the texture element extreme points are the same, the problem is further simplified. Now, any side of the texture element will match seamlessly with any other side, so the texture element may be placed on the triangle in any orientation. Simply assign coordinates to one vertex randomly and then assign the other coordinates in turn, walking in the same direction around the texture element and triangle.

Clearly this simplification of the texture mapping process is not suitable for the creation of realistic textures on familiar objects, but it suffices for the creation of small scale, regular texture patterns on unfamiliar surfaces for the purpose of enhancing shape perception.

Triangular textures were created using *Pdbq*, a special-purpose language for molecular visualization developed by Palmer [Palmer92]. Pdbq is a C-like interpreted language which provides basic and geometric data types, control structures, overloaded operators, streams to support both simple and geometric file I/O, and built-in functions to perform standard, molecular, and geometric tasks. Simple, symmetric triangular textures [Figure 1] can be generated by such methods as:

```
{spots at texel center -- Fig 1a}
for all points (x,y)
   opacity(x,y) = 1 - dist((x,y), centroid)
```

or:

```
{spots at vertices (v1,v2,v3) -- Fig 1b}
for all points (x,y)
   opacity(x,y) =
      1 - min(dist((x,y),v1),
         dist((x,y),v2), dist((x,y),v3)
```

or:

```
{ring texture: RAD sets radius -- Fig 1c }
for all points (x,y)
   opacity(x,y) =
      1-MAX(ABS(RAD-min(dist((x,y),v1),
      dist((x,y),v2),dist((x,y),v3)
```

Designing textures which modulate opacity offer specific challenges. An effectiveß texture will have both enough transparent parts for the objects behind to be seen clearly and enough opaque parts for the surface on which it lies to be clearly perceived. Unfortunately, these requirements are directly at odds with one another. The right balance must be struck on a case-by-case basis, depending on clarity of the inside objects, the irregularity of the outside surface, and the relative importance of the various objects.

5. Tessellation

Obviously, the most regular texture patterns will be formed over surfaces with uniformly sized triangles. On such surfaces, texture elements are replicated at a constant size across the surface. Additionally, since the texture element is defined on an equilateral triangle, the textured surface will appear most regular when triangles of the surface are also equilateral. Clearly, an arbitrary surface will not necessarily be comprised of uniformly sized equilateral triangles. On

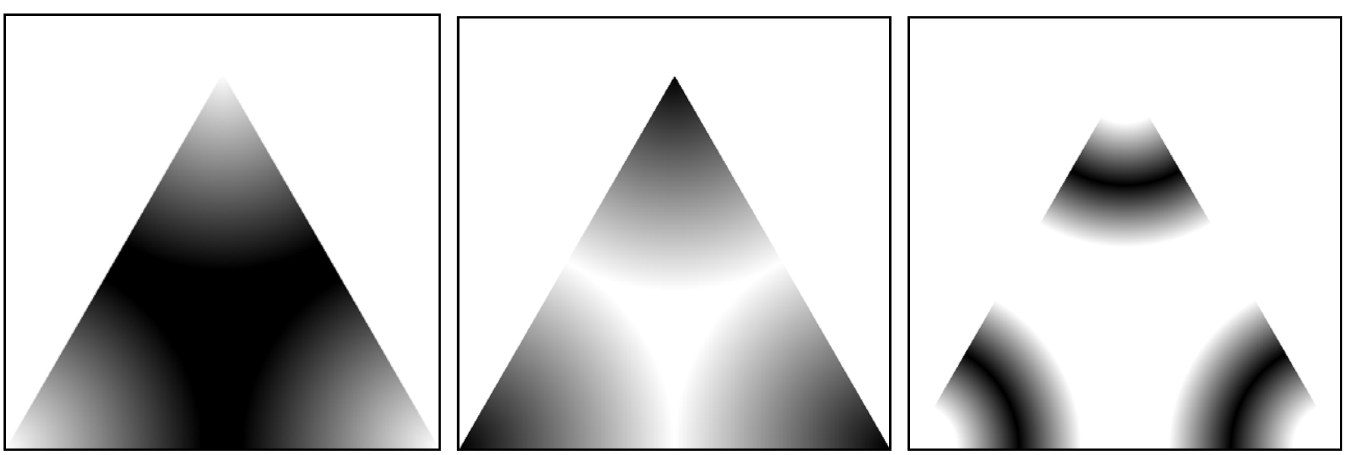

Figure 1. Three simple triangular textures with opaque area a) at element centroid, b) at vertices, and c) in rings around vertices. Dark areas represent high opacity.

irregular surfaces, triangles which are unusually large, small, or skinny result in disruptions of the texture pattern. Figure 2a shows the initial tessellation resulting from the computation of the solvent-accessible surface of a molecule. Figure 2b shows the result of texturing this surface with a triangular texture. Disruptions of the texture regularity are especially visible in the lower left part of the molecule. The textures used in Figure 2 are completely opaque to show the effects of the tessellation more clearly.

Texture pattern regularity can be improved by a preprocessing step which regularizes the polygonal tessellation, producing triangles which are more uniformly sized and closer to equilateral. This re-tiling was performed using an automatic method developed by Greg Turk [Turk92]. First, a new set of points is chosen at random locations on the surface. These points will later become the vertices of the regularized tessellation. Then, a relaxation procedure is applied to move points away from neighboring points by way of a simulated repulsion force. The result of this process is a set of candidate points uniformly distributed across the polygonal surface. Next, each polygon of the original triangulation is replaced by a set of triangles which uses the original polygon vertices and the candidate vertices which lie in the polygon to exactly tile the polygon. This mutual tessellation preserves the connectedness of the original surface. Finally, the original vertices are removed from the mutual tessellation, resulting in a polygonal surface containing only the new vertices. Figure 3 shows the result of regularizing the molecular

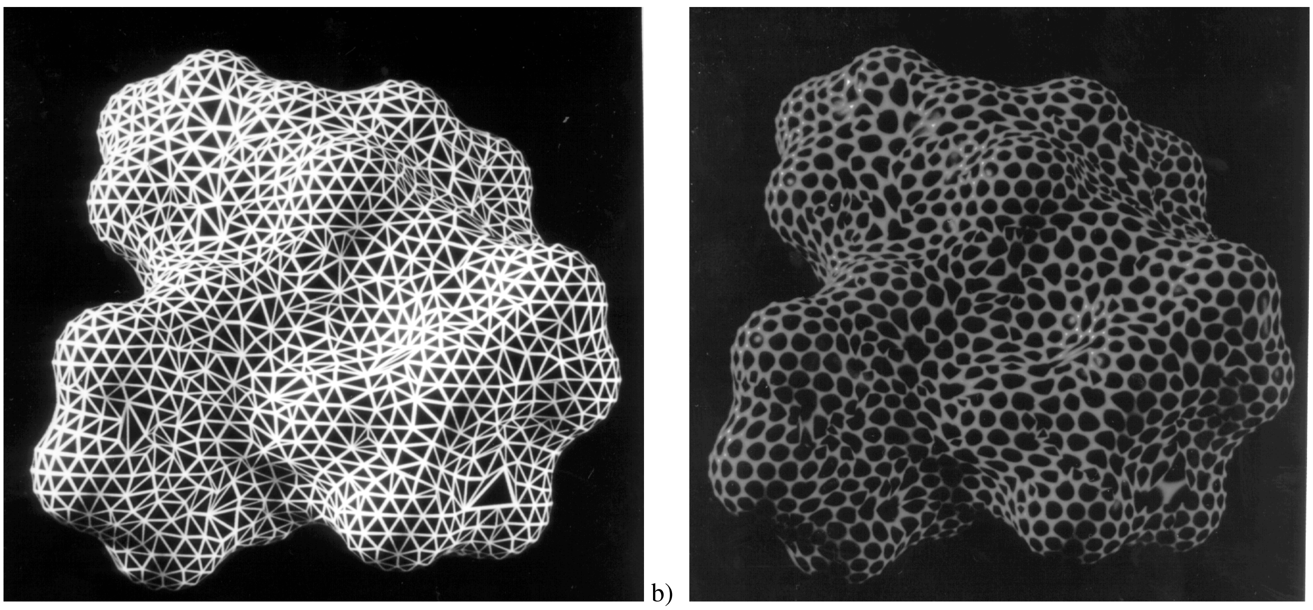

a) b)

Figure 2. a) Initial tessellation of solvent-accessible surface and b) result of applying triangular texture to that surface.

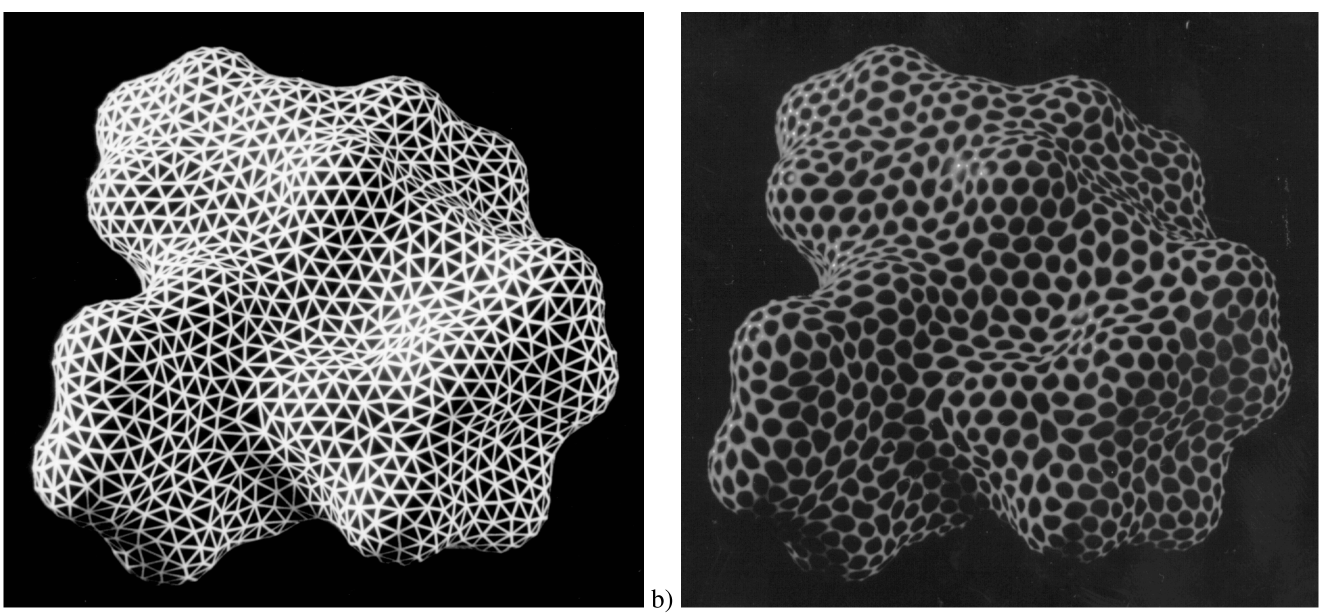

a) b)

Figure 3. a) Regularized tessellation of solvent-accessible surface and b) result of applying triangular texture to the new surface.

a) 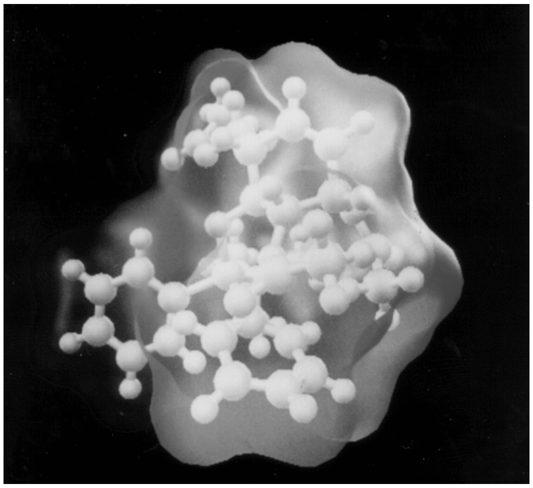 b) 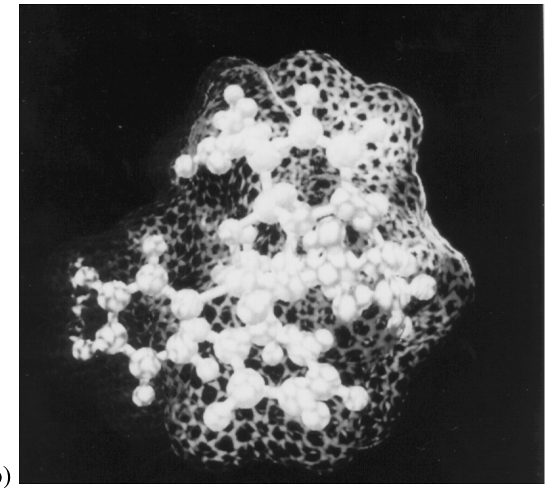

Figure 4. Comparison of a) transparent surface to b) opacity-modulating surface for molecular application. Color represents an additional variable (in this case, simply position in x) evaluated at the solvent-accessible surface. See color plates.

surface shown in Figure 2. Notice that triangles are distributed more evenly, resulting in a more regular texture pattern.

In fact, even just the hidden surface view of the new tessellation (Figure 3a) shows the molecule shape clearly, since the polygon boundaries form a relatively small scale, regular pattern. This image was rendered by the same method as the textured surfaces, in this case a triangular texture with high value near the boundaries of the texture element was used. This pattern, however, would not work well for an opacity-modulating texture, since the opaque and transparent areas are not properly balanced.

Although these same methods can be used to create polygonal models at various levels of detail, the molecular surfaces used in the figures were re-tiled with approximately the same number of polygons as the original tessellation, creating a more regular tessellation at the same level of detail.

6. Results

Details of the shape of the molecular surface are visible in the textured image. For instance, the vertically compressed texture pattern near the center of the image in Figure 3b indicates a “shelf” in the molecular surface. A similar feature is visible near the indentation on the right side of the molecule. In an untextured surface, the visibility of these features would be dependent on the position of lights in the scene. For many lighting configurations, these features would be extremely difficult to see.

Figure 4 shows the solvent-accessible surface of the same molecule represented with a transparent surface (Figure 4a) and an opacity-modulating textured surface (Figure 4b). In both cases the ball-and-stick representation of the molecule’s atoms and bonds can clearly be seen within the surface. The textured image gives better clues to the shape of the surface, particularly in places where the surface is nearly orthogonal to the view direction. Both images require close examination to understand the shape of such a complex, unfamiliar object, but at some places on the surface of the transparent object no amount of examination seems to reveal structural details. In both images color has been used to represent the values of another variable at points on the surface. While the overall distribution of color values is more striking in the transparent surface, in areas where elements of the ball-and-stick representation underlie the surface, the color is very difficult to discern. In the textured surface, color values are apparent at opaque parts of the surface. However, if the opaque bands are narrow, judging color values may require careful examination. While color appearance of transparent surfaces is always changed by the color of underlying surfaces, the color appearance of opaque parts of an opacity-modulating surface is only influenced by simultaneous contrast effects with nearby areas, including those visible through holes in the surface.

Statistical surfaces textured with opacity-modulating patterns are also more comprehensible than those rendered as transparent surfaces. Figure 5 shows a statistical surface of ozone concentration represented using both methods. In both images, ozone level is mapped redundantly to height and color. The surface in Figure 5b is textured with a rectangular opacity-modulating texture, but the same general principles as for triangular textures apply. Both techniques show the map outline lying below the surface, but texturing conveys more information about surface shape. In the textured image, the slight valley between the two ridges in the foreground is more visible. In the transparent image, the ridges themselves are visible primarily from color changes and not from shape cues. Similarly, subtle undulations of the far left part of the surface are clearly visible in the textured image but appear to be a formless jumble in the transparent image.

The shape of multiple irregular molecular surfaces is virtually incomprehensible when displayed using transparent surfaces [Figure 6a]. In these images, the solvent-accessible surfaces for solvent atoms of two different radii are displayed along with the ball-and-stick representation. Surface structure where the screen projections of the different surfaces overlap is particularly difficult to discern. Display of the same surfaces using opacity-modulating textures is somewhat clearer [Figure 6b], especially in places where surface orientation causes substantial distortions of the texture gradient. Still, a clear understanding of the details of surface shape is difficult to glean from either image. When the image is animated, however, the improvement in clarity resulting from the textured surface is striking. The more abundant surface features provided by the repeating texture facilitate perception of the surface shape. Relative movement of these small features creates a strong sensation of depth relationships, improving surface perception.

7. Extensions

Stereo pairs of textured irregular surfaces should offer many of the benefits of rotation of the surfaces at a much lower computational cost. Just as moving textured surfaces offer compelling shape cues not present in either static views or moving transparent surfaces, so should stereo views of textured surfaces. The richness of textural detail provides abundant features for matching locations between the two views. Since it is the positional disparity of features between the two views which creates the sensation of depth, the presence of adequate textural features ensures a strong

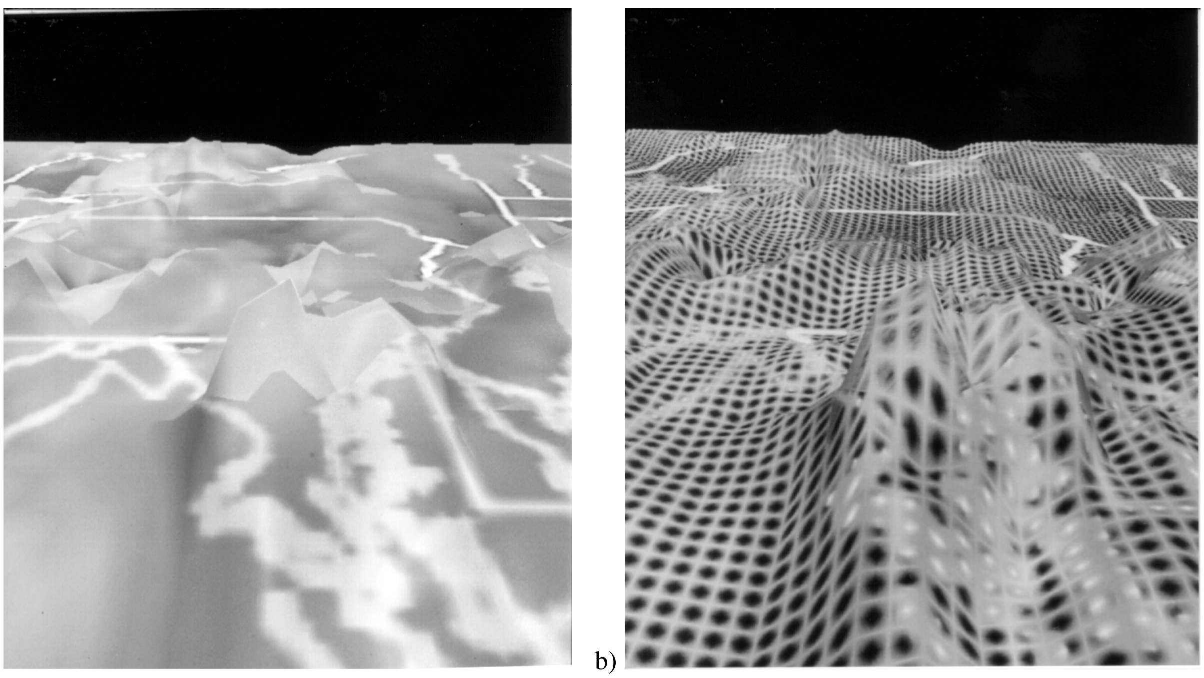

a) b)

Figure 5. Comparison of a) transparent and b) textured statistical surface. Ozone concentrations over the Eastern seaboard of the US are redundantly mapped to height and color in both images. See color plates.

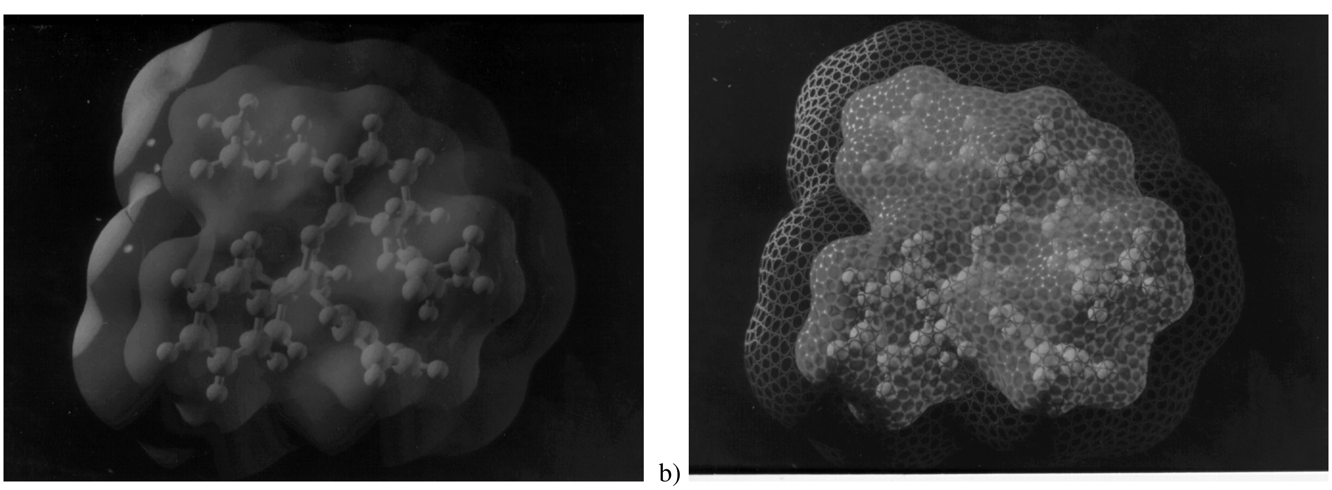

a) b)

Figure 6. Comparison of a) multiple transparent surfaces to b) multiple opacity-modulating surfaces. See color plates.

perception of depth from positional disparity, just as it ensures a strong perception of depth from relative motion. This expectation is supported by the fact that the mechanisms of stereo perception are similar to those of motion perception, and by and large the two are performed by the same channels in the human visual system.

Another interesting question which remains is whether the texture pattern can be used to carry information of its own without disrupting its role in perception improvement. For instance, type of texture element could encode some classifier on the surface, such as tissue type or land use. As long as texture elements joined seamlessly at polygon boundaries, either because classification of adjoining polygons was guaranteed to match or because all texture elements had identical values along their boundaries, no complications would be introduced to the texture mapping process. Alternatively, texture density or opacity drop-off rate could represent confidence in the data. Such a technique would require that the texture be procedurally generated, unless the number of distinct values represented by texture characteristics was fairly small. Otherwise, a potentially prohibitive number of different texture elements would need to be generated and stored in order to represent the range of values (or combinations of values).

Acknowledgments

This work performed at the North Carolina Supercomputer Center and the US EPA Visualization Center. Greg Turk contributed ideas, encouragement, and the source code for his polygonal re-tiling methods. Tom Palmer provided valuable discussions and the molecular visualization language Pdbq. The manufactured carbohydrate named 'Wilma' was provided by Mark Zottola. Ozone concentration data came from the US EPA Atmospheric Research and Exposure Assessment Laboratory.

References

[Cumming93] Cumming, Bruce, Elizabeth Johnston, and Andrew Parker, Effects of Different Texture Cues on Curved Surfaces Viewed Stereoscopically. *Vision Research*, vol. 33, no 5/6, 1993, pp. 827-838.

[Cutting84] Cutting, James, and Robert Millard, Three Gradients and the Perception and Flat and Curved Surfaces. *Journal of Experimental Psychology: General*, vol. 113, no. 2, 1984, pp. 198-216.

[Dooley90] Dooley, Debra, an Michael Cohen, Automatic Illustration of 3D Geometric Models: Surfaces. *Proceedings of Visualization '90*, October 1990, pp. 307-314.

[Duncan95] Duncan, Bruce, and Arthur Olson, Texture mapping parametric molecular surfaces. *Journal of Molecular Graphics*, vol. 13, no. 4, August 1995, pp. 258-264.

[Gibson50] Gibson, James, *The Perception of the Visual World,* Houghton Mifflin, 1950.

[Interrante95] Interrante, Victoria, Enhancing Transparent Skin Surfaces with Ridge and Valley Lines. *Proceedings of Visualization '95*, October 1995, pp. 52-59.

[Levoy90] Levoy, Marc, Henry Fuchs, Stephen Pizer, Julian Rosenman, Edward Chaney, George Sherouse, Victoria Interrante, and Jeffrey Kiel, Volume Rendering in Radiation Treatment Planning. *First Conference on Visualization in Biomedical Computing*, 1990, pp. 4-10.

[Palmer92] Palmer, Thomas C., A Language for Molecular Visualization. *IEEE Computer Graphics and Applications*, vol. 12, no. 3, May 1992, pp. 23-32.

[Schweitzer83] Schweitzer, Dino, Artificial Texturing: An Aid to Surface Visualization. In Computer Graphics (SIGGRAPH 83 Proceedings], vol. 17, no. 3, July 1983, pp. 23-27.

[Teschner94] Teschner, Michael, Christian Henn, Horst Vollhardt, Stefan Reiling, and Jurgen Brickman, Texture mapping: A new tool for molecular graphics. *Journal of Molecular Graphics*, vol. 12, no. 2, June 194, pp. 98-105.

[Todd87] Todd, James, and Robin Akerstrom, Perception of Three-Dimensional Form from Patterns of Optical Texture. *Journal of Experimental Psychology: Human Perception and Performance,* vol. 13, no. 2, 1987, pp. 242-255.

[Turk92] Turk, Greg, Re-Tiling Polygonal Surfaces. In *Computer Graphics (SIGGRAPH 92 Proceedings)*, vol. 26, no. 2, August 1992, pp. 55-64.

[Turk91] Turk, Greg, Generating Textures on Arbitrary Surfaces Using Reaction-Diffusion. In *Computer Graphics (SIGGRAPH 91 Proceedings)*, vol. 25, no. 4, July 1991, pp. 289-298.

[vanDamme94] van Damme, Wim, *Active Vision: exploration of three-dimensional structure*. Ph.D. thesis, University of Utrecht, 1994, ISBN 90-393-0803-0.

[Wejchert92] Wejchert, Jakub, Textures and Equation Visualization. *IEEE Computer Graphics and Applications,* vol. 12, no. 6, November 1992, pp. 10-12.

Figure 4. Comparison of a) transparent surface to b) opacity-modulating surface for molecular application. Color represents an additional variable (in this case, simply position in x) evaluated at the solvent-accessible surface.

Figure 5. Comparison of a) transparent and b) textured statistical surface. Ozone concentrations over the Eastern seaboard of the US are mapped to height and color in both images.

Figure 6. Comparison of a) multiple transparent surfaces to b) multiple opacity-modulating surfaces.