Enhancing Transparent Skin Surfaces with Ridge and Valley Lines*

Victoria Interrante, Henry Fuchs, Stephen Pizer
Department of Computer Science
University of North Carolina at Chapel Hill

Abstract

There are many applications that can benefit from the simultaneous display of multiple layers of data. The objective in these cases is to render the layered surfaces in a such way that the outer structures can be seen and seen through at the same time. This paper focuses on the particular application of radiation therapy treatment planning, in which physicians need to understand the three-dimensional distribution of radiation dose in the context of patient anatomy.

We describe a promising technique for communicating the shape and position of the transparent skin surface while at the same time minimally occluding underlying isointensity dose surfaces and anatomical objects: adding a sparse, opaque texture comprised of a small set of carefully-chosen lines. We explain the perceptual motivation for explicitly drawing ridge and valley curves on a transparent surface, describe straightforward mathematical techniques for detecting and rendering these lines, and propose a small number of reasonably effective methods for selectively emphasizing the most perceptually relevant lines in the display.

1: Introduction

The goal of radiation therapy is to maximally irradiate the target volume, which includes the tumor along with a slight margin, while delivering minimum radiation to healthy tissue and completely sparing certain sensitive anatomical structures. In designing a treatment, the physician defines the number, direction, intensity and shape of the radiation beams to be used and then computes the dose distribution that would result from applying this beam configuration to the patient, based on his CT scans. To evaluate this proposed treatment, the physician needs to visualize the three-dimensional distribution of radiation dose superimposed over the patient anatomy. This typically requires the display of multiple layers of information, including the target, sensitive anatomical structures, radiation dose and anatomical context.

There are several advantages of using transparency to display everything at once in three dimensions. The first is simultaneity: displaying all relevant external and internal anatomical objects together with one or more isointensity surfaces of radiation dose allows better comprehension of the complex spatial relationships between the various structures. Psychophysical experiments have shown that we are measurably better at estimating the distances between two objects when the objects are shown together in a single image than when we are forced to remember the location of one of the objects from a previous image [4]. The second advantage is completeness: important features of complicated 3D shapes can be overlooked when the data is viewed only as a sequence of 2D slices and then mentally reconstructed [6]; 3D display allows maximum comprehension of 3D form. Finally, there is the advantage of efficiency: if the physician is able to obtain more of the information he needs about the data from a smaller number of displayed images, the time required to evaluate a particular plan can be reduced and a greater number of different treatment plans can potentially be explored [17].

The primary disadvantage of using transparency to display geometrically complex superimposed surfaces is that, in computer graphics as in real life, it is often difficult to adequately perceive the three-dimensional shape and location of a transparent surface except at its silhouette edges. The effects of refraction often seem to do more to confuse the perception of underlying information than to provide cues to the shape of the transparent surface, and the ordinary shape and depth cues of shading and occlusion are only minimally present on transparent surfaces, making it particularly difficult to appreciate their position and orientation. Figure 1, a photograph of a typical transparent surface superimposed over a familiar background, illustrates some of the difficulties we encounter in understanding the shape and depth of layered transparent surfaces in real life. It is fairly clear that we need to add something to the transparent surface to more explicitly define it. A sparse, opaque texture could help provide useful cues to the surface's shape and relative location in depth while at the same time allowing an unimpeded view of most of the underlying information. Because the data will typically need to be examined from multiple points of view, it will be important to ensure that this texture remains stable under conditions of changing surface orientation.

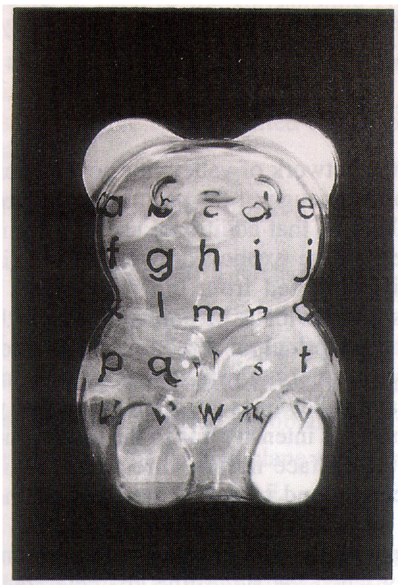

Figure 1: photograph of a transparent object

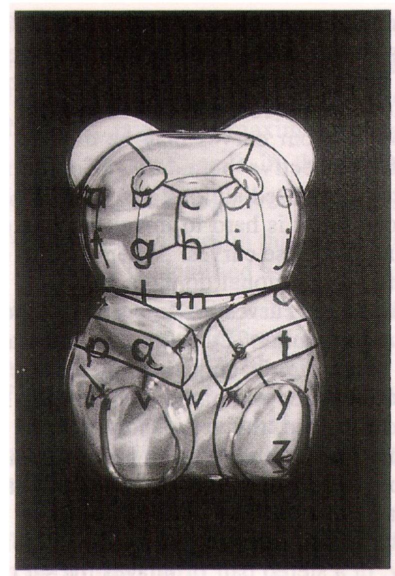

Figure 2: photograph of a transparent object with ridge and valley line texture

Inspired by the ability of gifted artists to define a figure with just a few strokes, and based on a reasonably thorough investigation into the psychology of form perception, we have designed and implemented a method for generating a stable, perceptually intuitive, 3D “line drawing” of a surface based on local measures of its differential geometry. This concept is illustrated in figure 2, a photograph that shows a (fairly well-informed) estimation of the major ridge and valley lines superimposed in 3D on the transparent surface of figure 1 using drafting tape.

2: Related Work

Dooley and Cohen [5] use screen-space texture patterns to clarify the depth order of layered transparent surfaces. Levoy *et al* [12] use solid grid textures on layered transparent surfaces in volume data to enhance the communication of shape and position information. Pearson and Robinson [15] use “line drawings” derived from luminance valleys in 2D images for the communication of facial and gesture information at very low data rates. Ponce and Brady [16] compute a “surface primal sketch” from lines of C^0 and C^1 discontinuity on height surfaces defined by range maps. Koenderink [9] provides a clear mathematical definition of ridge lines, which he includes in his list of singular surface features, and mentions applications for which it might be useful to use ridge lines to divide a surface into patches. Bookstein and Cutting [2] compute ridge lines on surfaces in 3D craniofacial data and use them as landmark curves for morphometrical analyses. Monga *et al* [14] compute ridge lines on isointensity surfaces in 3D volume data and use them for data registration [19] and automatic atlas generation [18]. The approach taken in this work to detect and display ridge and valley lines, although basically very similar to the method of Monga *et al*, is perhaps better characterized as a fairly straightforward implementation of the ridge definition provided by Koenderink. The principal difference between our method and that of the INRIA group is that, because we are concerned more with the display of perceptually relevant features than with locating exact ridge lines, we forego the computation of third derivatives in favor of a very stable and simple approximation that tests for the presence of a local curvature maximum in a subvoxel region.

3: Defining a 3D line drawing

Our goal in generating a sparse, opaque texture for transparent surfaces is to come up with a very small set of meaningful curves that describe the data well. We want to capture the essence of the outermost surface shape in a few, well-chosen lines and leave the underlying data clearly visible through the remaining portions of the transparency.

3.1: The role of silhouette and contour curves

Silhouette and contour curves are the 2D projection of points on the 3D surface where the direction of the surface normal is orthogonal to the line of sight. Silhouette curves form a closed outline around the projection; contour curves may be disjoint and can fall within the projective boundary. Contours (including silhouettes) are used ubiquitously in 2D line art and illustration, and it is hard to imagine a line drawing that doesn't include these curves. When we turn our attention to the three-dimensional domain, however, we find that although contour curves are important shape descriptors [8] for a particular 2D, static, monocular view, the benefits of explicitly emphasizing these curves under conditions of stereo and motion are limited. As an object is rotated in space, the set of points on its surface that map to the contour changes with each reorientation. In an animation, or under conditions of dynamic object manipulation, explicitly marked contour curves will appear to crawl over the three-dimensional surface. Similarly, in a stereo pair of images the set of points that map to the contour in the view from one eye will be different from the set of points that map to the contour in the view from the other eye. Because our tendency when viewing a stereo pair is to establish a correspondence between similar points in the two images, drawing the same surface points differently in the two views can confuse the perception of the surface data. For these reasons, we have rejected the use of silhouette and contour lines as a sparse surface texture in favor of another sparse set of descriptive lines that remain fixed on the surface under dynamic viewing conditions. While they are not selectively emphasized in our renderings, silhouette edges do remain clearly visible in the display as the locus of points between the figure and the background.

3.2: The perceptual relevance of valley lines

As defined by Koenderink [9], valley lines are the locus of points on a surface at which the normal curvature assumes a local minimum in the principal direction associated with the largest, negative curvature. Explicitly drawing the set of valley lines on a surface makes sense perceptually for several reasons.

On ordinary matte surfaces, the effects of self-shadowing tend to leave valley regions less brightly lit than surrounding surface areas. Artists' tendency to draw black lines along valley creases is in part a reflection of this shading effect. Although transparent surfaces don't exhibit the same shape from shading characteristics as opaque, matte surfaces, we can help provide an intuitively meaningful surface description by taking advantage of some of these general shading cues, making them explicit in an orientation-invariant way.

Psychophysical studies [3] have verified theories [7][1] that people tend to perceive objects as naturally partitioning along their valley lines. It follows that by making these lines explicit we are, in effect, if not drawing in what people already think they see, at least illustratively complementing a naturally-existing perceptual process.

3.3: The role of ridge lines

When the network of valley lines over a surface does not incorporate enough of the prominent surface features, it may be the case that additional useful information can be conveyed by the explicit drawing of the ridge lines, possibly differentiated from the valley lines by color. Ridge lines are the locus of points at which the normal curvature in the principal direction associated with the largest, positive curvature assumes a local maximum. Both ridge and valley lines are often associated with sharp changes in surface intensity due to the locally rapid rate of change of the surface normal direction. While valleys generally correspond to darker portions of the surface, ridges are more likely to reflect a specular highlight. Drawing lines along ridges seems to make the most sense perceptually when the ridge line characterizes an easily recognizable surface shape feature, such as the bridge of the nose, or marks a shading discontinuity- which occur more commonly on manmade surfaces than in biological data.

4: Implementation

Both ridge and valley lines correspond to geometric features of the surface, and as such they can be computed automatically from local measures of the surface's differential geometry. Our rendering approach is based on an integration of the marching cubes isosurface detection algorithm developed by Lorensen and Cline [13] and the ray-casting method of Levoy [11], whose volume rendering software platform forms the base upon which most of this display code has been built.

4.1: Locating the surface in the volume

The first step, for each pixel, is to locate the point at which the viewing ray through that pixel intersects the transparent surface, defined as an isointensity surface in the volume. We do this by examining data values at the eight corner voxels of successive data cells pierced by the ray as it traverses the volume. If an isosurface crossing is detected, the surface triangulation in the cell is computed, the ray is intersected with these polygons to define the exact ray/surface intersection point in 3D, and then the temporary polygonal representation is discarded. Only the first intersection with each transparent surface is displayed; subsequent intersections of a ray with the same isointensity surface are ignored, to improve image clarity.

4.2: Computing smoothed surface normals

We define a smoothed surface normal at the ray/surface intersection point by interpolating from grey-level gradients at each of the surrounding eight voxels, which were precomputed using a Gaussian-weighted filter. The use of Gaussian smoothing and floating point normals turns out to be essential for reducing directional artifacts that arise due to the limited (8-bit) precision of the volume data.

4.3: Defining an orthogonal frame

Once we have located a surface point $P_{x,y,z}$ and computed its normal direction, we can define an orthogonal frame ($\vec{e}_1, \vec{e}_2, \vec{e}_3$) at that position, in which $\vec{e}_3$ points in the surface normal direction and $\vec{e}_1$ and $\vec{e}_2$ span the tangent plane. In practice we obtain $\vec{e}_1$ by choosing an arbitrary direction the tangent plane and then get $\vec{e}_2$ by taking the cross-product of $\vec{e}_1$ and $\vec{e}_3$.

4.4: Computing principal directions, maximum and minimum normal curvatures

From the orthogonal frame, we can determine the Second Fundamental Form- a matrix $A = \begin{bmatrix} \omega_1^{13} & \omega_1^{23} \\ \omega_2^{13} & \omega_2^{23} \end{bmatrix}$ that describes the local surface shape in terms of changes in the direction of the surface normal in the local area. The coefficients ω_j^{i3} specify the component in the $\vec{e}_i$ direction of the rate at which the normal tips as you move across the surface in the $\vec{e}_j$ direction. Koenderink describes the terms as representing "nosedives" when $i = j$ and "twists" when $i \neq j$. We compute ω_j^{i3} by taking the dot product of $\vec{e}_i$ and the first derivative of the gradient in the $\vec{e}_j$ direction. The determinant of the Second Fundamental Form matrix is the Gaussian curvature. If we rotate the frame in the tangent plane so that the twist terms ω_2^{13} and ω_1^{23} disappear, diagonalizing A to get $D = \begin{bmatrix} k_1 & 0 \\ 0 & k_2 \end{bmatrix}$ and $P = \begin{bmatrix} \mathrm{v}_{1_u} & \mathrm{v}_{2_u} \\ \mathrm{v}_{1_v} & \mathrm{v}_{2_v} \end{bmatrix}$ ($A = PDP^{-1}$ and $|k_1| > |k_2|$), then the principal directions- the two orthogonal directions in which the normal curvatures assume a maximum and a minimum value- are given by $\vec{e}_i' = \mathrm{v}_{i_u}\vec{e}_1 + \mathrm{v}_{i_v}\vec{e}_2$ and the principal curvatures are specified by k_1 and k_2.

4.5: Determining whether a point is on or near a ridge or valley

After we have computed the principal directions, we can label the point $P_{x,y,z}$ as being "near a ridge" if $k_1 > 0$ and k_1 is a local maximum of normal curvature in the first principal direction at the point $P_{x,y,z}$ or as being "near a valley" if $k_1 < 0$ and k_1 is a local minimum of normal curvature in the first principal direction. We determine whether k_1 is a local maximum or minimum by stepping out a short distance in the volume from $P_{x,y,z}$ in the positive and negative $\vec{e}_1'$ directions and comparing the normal curvatures in this direction at these points to k_1. If $P_{x,y,z}$ is identified as being near a ridge or valley, it is assigned an additional amount of opacity and in some cases a slightly different color in order to identify it more prominently in the image. Volume rendering of other data objects proceeds in the usual manner, and the color and opacity values of all surfaces encountered along the ray are composited in front-to-back order to produce a final pixel value for display.

5: Improving the display of ridge and valley information

If all of the ridge and valley points that are identified on the transparent surface by the preceding definition are displayed opaquely, the result, illustrated in figure 3a, is somewhat less than satisfactory. It is clear that additional steps must be taken to selectively emphasize the more prominent ridge and valley regions while de-emphasizing the others.

5.1: Defining opacity as a function of normal curvature in the first principal direction

The first step we take is to define the opacity of the ridge or valley points as a function of the relative magnitude of the normal curvature in the first principal direction. This reinforces the presence of the texture at more sharply curved regions while minimizing it in flatter places. While it appears that this approach requires us to first obtain an estimate of the global maximum and minimum k_1 values over all ridge and valley points in the image, in practice we find that with some experience suitable maximum and minimum k_1 values are fairly easy to estimate and that in any event the results, illustrated in figure 3b, are not overly sensitive to the specific values chosen for $k_{\max}$ and $k_{\min}$.

5.2: Adding curvature-based thresholding

Additional improvements in the effectiveness of the display, illustrated in figure 3c, are obtained by defining threshold curvature values; a ridge or valley point is displayed with additional opacity only if the magnitude of the maximum principal curvature at that point exceeds the specified cutoff.

3a: all detected ridge and valley points

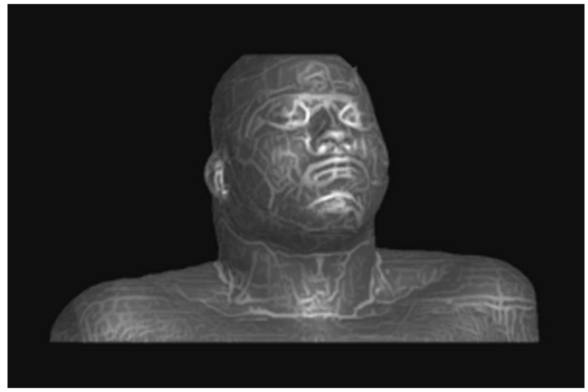

3b: opacity as a function of principal curvature

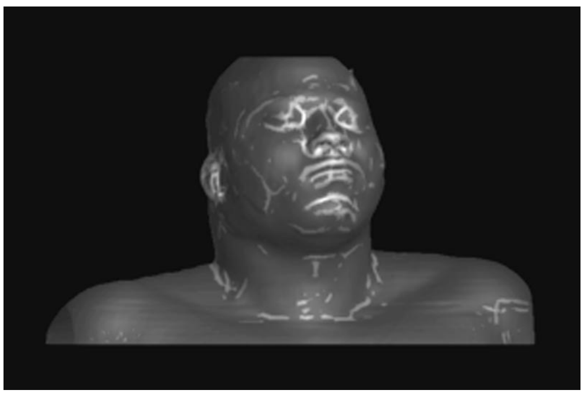

3c: curvature-based opacity threshold

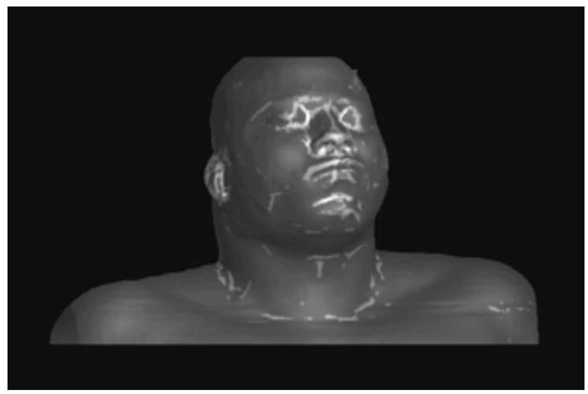

3d: approximate feature width criterion used

Figure 3: Improving the display of ridge and valley information

5.3: Applying an approximate minimum width criterion

In certain cases, we find that very deep but extremely narrow spurious ridge or valley creases remain visible after the aforementioned techniques are applied. In order to identify these points and eliminate them from the texture, we need to look at a slightly larger local region of the surface. One approach that has shown some limited promise is to step away on either side of $P_{x,y,z}$ in the first principal direction, which in general is not orthogonal to the ridge direction but can serve our purposes nevertheless, and compare the angles between $\vec{e}_3$ and the gradients at these successive points. If we find that the approximated surface normals begin to realign before a specified minimum distance has been traversed, we conclude that the ridge or valley feature is narrower than desired and no additional opacity is assigned to the point. The results we obtain using this approximate minimum width criterion in addition to the strictly local curvature criteria are illustrated in figure 3d.

6: Results

Figure 4a shows a radiation therapy treatment plan for a patient with cancer of the nasopharynx, a region at the top of the throat and just behind the nose. This is actual clinical data, provided by our colleagues at UNC hospitals. The outer skin surface has been rendered semitransparently, and through this surface we can see the opaque treatment region enclosed by a semitransparent isointensity surface of radiation dose. In figure 4b, the same data has been rendered with ridge and valley line texture added to the skin surface. A standard solid grid texture is shown in figure 4c for purposes of comparison.

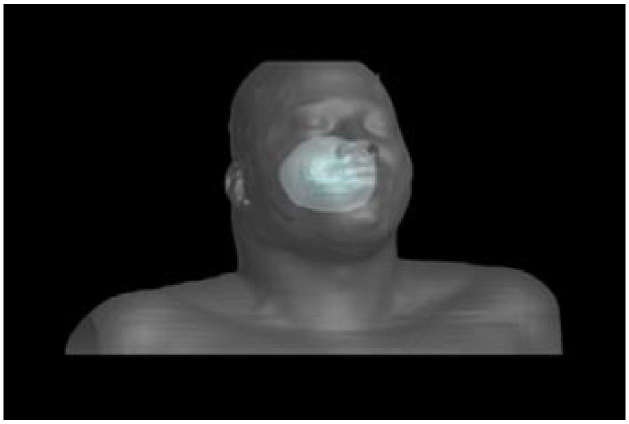

4a: untextured skin surface

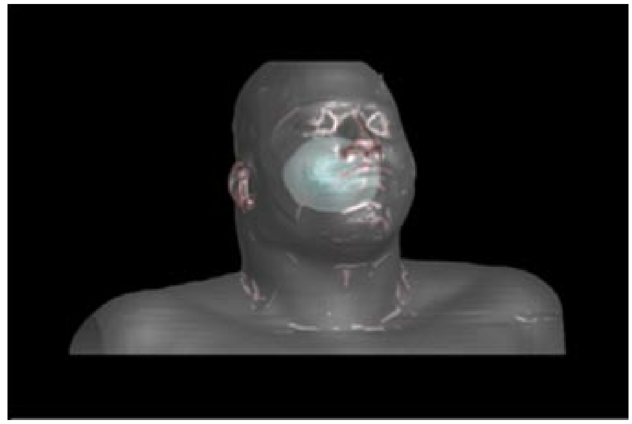

4b: skin with ridge and valley line texture

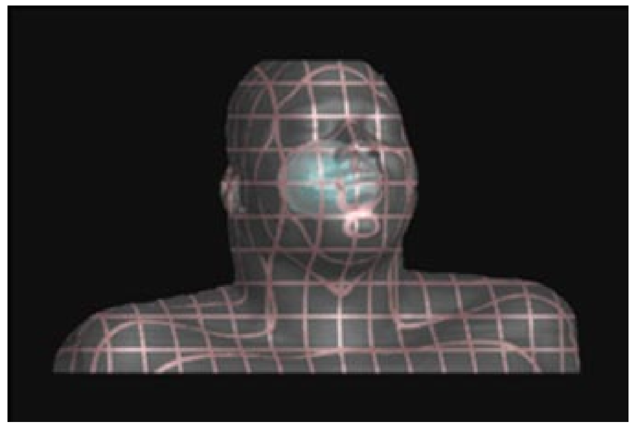

4c: skin with solid grid texture

Figure 4: Radiation therapy treatment plan for cancer of the nasopharynx

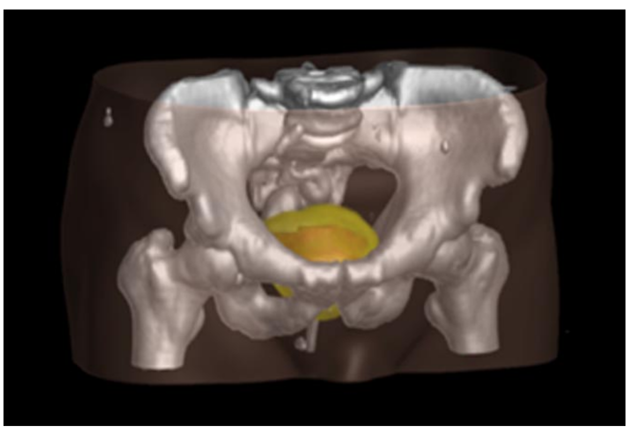

5a: untextured skin surface

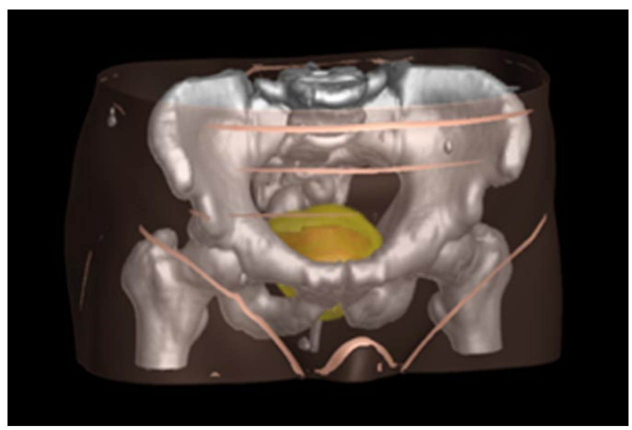

5b: skin with ridge and valley line texture

Figure 5: Radiation therapy treatment plan for cancer of the prostate

In figure 5a we show a treatment plan for a patient with prostate cancer. A semitransparent skin surface is displayed superimposed over the bones of the pelvis, within which we can see the opaque treatment region and an enclosing, closely matched isointensity surface of radiation dose. Figure 5b shows the same dataset with an opaque valley line texture added to the skin surface.

7: Future Work

We have not explored in great depth the possibility of applying different colors to the ridge and valley lines. For the illustrations in this text, we have used a very subtly darker shade of the skin color for the valley lines and a subtly lighter shade for the ridges. While our implementation allows arbitrary colors to be applied to each of the line types, it is our impression (based more on intuition than on experience) that there would be little value in explicitly labeling this distinction through the use of different hues; to avoid overloading these already complex visualizations with too much extraneous information, it might be better to reserve the use of color for differentiating between the various underlying dose surfaces and anatomical structures.

Although the texturing technique described in this paper appears promising, we cannot yet make any substantive claims about its utility in clinical practice. It has been repeatedly shown that results from preference studies are poor predictors of performance benefits [10]. To obtain a meaningful measure of the efficacy of ridge and valley line textures, we would need to conduct controlled user experiments comparing specific, task-related performance using textured and untextured data; certain studies in this category are in progress.

8: Conclusions

Adding a sparse, opaque texture to a transparent surface can help make its location in space much more explicit, providing additional occlusion cues and possibly enabling a better estimation of relative depth from motion. We provide examples suggesting that better results can be perhaps be obtained when the occluding matter we add somehow “makes sense” intuitively. For certain familiar surfaces that are easily characterized by a few recognizable shape-based landmarks, it appears that ridge and valley lines can serve as an excellent meaningful sparse texture. A primary advantage of using ridge and valley lines rather than other see-through textures is that, in addition to providing areas of opacity, the lines themselves carry both geometrical and perceptually relevant information.

In the visualization of radiation therapy treatment planning data, displaying valley and ridge lines on a transparent skin surface may enable us to better communicate the existence, form, and location in depth of this surface relative to the underlying target and dose structures while only minimally occluding them. Presenting a comprehensible skin surface can be useful for providing an overall context (scale and orientation) for the target and dose, and for making explicit some of the sensitive soft tissue areas that should be avoided by the radiation beams. Representing the skin surface by a few, well-chosen lines can help us to communicate this essential information without compromising the visibility of important underlying data; ridge and valley lines appear particularly well-suited for this task.

Acknowledgments

This work was supported by NIH grant # PO1 CA47982. The authors would like to thank Fred Brooks, Christina Burbeck and Julian Rosenman for their helpful contributions. We are grateful to Marc Levoy for making available his volume rendering source code, and to Jim Chung for providing the implementation of the marching cubes algorithm that we use for the surface triangulation.

References

[1] Irving Biederman, “Recognition-by-Components: A Theory of Human Image Understanding”, *Psychological Review*, 1987, vol. 94, pp. 115-147.

[2] Fred Bookstein and Court Cutting, “A Proposal for the Apprehension of Curving Craniofacial Form in Three Dimensions”, Craniofacial Morphogenesis and Dysmorphogenesis, K. Vig and A. Burdi eds., 1988, pp. 127-140.

[3] Myron Braunstein, Donald Hoffman, Asad Saidpour, “Parts of Visual Objects: an Experimental Test of the Minima Rule”, *Perception*, 1989, vol. 18, pp. 817-826.

[4] Christina Burbeck and Yen Lee Yap, “Spatiotemporal Limitations in Bisection and Separation Discrimination”, *Vision Research*, 1990, vol. 30, no. 11, pp. 1573-1586.

[5] Debra Dooley and Michael Cohen, “Automatic Illustration of 3D Geometric Models: Surfaces”, *Visualization ‘90*, pp. 307-314.

[6] Elliot Fishman, Donna Magid, Derek Ney, Edward Chaney, Stephen Pizer, Julian Rosenman, David Levin, Michael Vannier, Janet Kuhlman and Douglas Robertson, “Three-Dimensional Imaging”, *Radiology*, 1991, vol. 181, pp. 321-337.

[7] Donald Hoffman and W. Richards, “Parts of Recognition”, *Cognition*, 1984, vol. 18, pp. 65-96.

[8] Jan Koenderink, “What Does the Occluding Contour Tell Us About Solid Shape?”, *Perception*, 1984, vol. 13, pp. 321-330.

[9] Jan Koenderink, Solid Shape, 1990.

[10] Harold Kundel, “Perception and Representation of

Medical Images", *Medical Imaging 1993. Image Processing*, SPIE vol. 1898, pp. 2-12.
[11] Marc Levoy, "Display of Surfaces from Volume Data", *IEEE Computer Graphics and Applications*, 1988, vol. 8, no. 3, pp. 29-37.
[12] Marc Levoy, 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*, May 1990, pp. 4-10.
[13] William Lorensen and Harvey Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm", *Computer Graphics*, 1987, vol. 21, no. 4, pp. 163-169.
[14] Olivier Monga, Serge Benayoun, Olivier Faugeras, "From Partial Derivatives of 3D Density Images to Ridge Lines", *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, June 1992, pp. 354-359.
[15] Don Pearson and John Robinson, "Visual Communication at Very Low Data Rates", *Proceedings of the IEEE*, 1985, vol. 73, no. 4, pp. 795-812.
[16] Jean Ponce and Michael Brady, "Toward a Surface Primal Sketch", *IEEE International Conference on Robotics and Automation*, March 1985, pp. 420-425.
[17] Julian Rosenman, Scott Sailer, George Sherouse, Edward Chaney and Joel Tepper, "Virtual Simulation: Initial Clinical Results", *Int. J Radiation Oncology Biol. Phys.*, 1991, vol. 20, pp. 843-851.
[18] Gerard Subsol, Jean-Philippe Thirion, Nicholas Ayache, "Steps Towards Automatic Building of Anatomical Atlases", *Visualization in Biomedical Computing*, Oct. 1994, pp. 435-446.
[19] J-P Thirion, O. Monga, S. Benayoun, A. Guezic, N. Ayache, "Automatic Registration of 3D Images Using Surface Curvature", *Mathematical Methods in Medical Imaging*, 1992, SPIE vol. 1768, pp. 206-216.