Three Dimensional Image Presentation Techniques in Medical Imaging

March 1987

Technical Report 87-009

Stephen M. Pizer and Henry Fuchs

The University of North Carolina at Chapel Hill
Department of Computer Science
New West Hall 035 A
Chapel Hill, N.C. 27514

To appear in Proc. Int. Symp. on Computer Assisted Radiology, Berlin, July 1-4, 1987 (Springer-Verlag) and Proc. The First European Conference on Imaging and Visual Documentation in Medicine, Amsterdam, May 24-27, 1987 (Elsevier)

Three Dimensional Image Presentation Techniques in Medical Imaging

STEPHEN M. PIZER*# and HENRY FUCHS*

Departments of Computer Science* and Radiology#
University of North Carolina
Chapel Hill, North Carolina, USA

Abstract

Medical images can be presented three-dimensionally by techniques that either calculate the effect of reflections from surfaces predefined from slices or project a three-space of luminosities computed from voxel intensities onto the visual receptors. Slice-based reflective displays are the most common type. Means of producing surface descriptions both via voxel sets and via slice contours will be reviewed. Advantages of and means of transparent display to allow the appreciation of the 3D relationships among objects will be set forth. Ways to produce additional depth cues by stereoscopy and the kinetic depth effect will be discussed, and the importance of interactive modification of viewpoint, clipping plane, displayed objects, etc. will be explained. A new device, UNC's Pixel-planes, for accomplishing this in real time will be illustrated.

Voxel intensity based display methods avoid the need for time-consuming predefinition of object surfaces and thus can allow exploration of 3D image data. Varifocal mirror hardware and fast computation of one or more projections based on object probabilities are two of the more important approaches.

While 3D display provides important information about 3D relationships, it cannot provide the kind of appreciation of subtle grey-scale changes that 2D display can. Methods that can combine these two kinds of information by superimposing 2D grey-scale slices on or in the context of 3D displays will be discussed.

Applications of these techniques for both diagnosis and radiotherapy planning will be used as illustrations and guides to the usefulness of these techniques with CT, MRI, and other 3D medical imaging modalities.

Introduction

Modern medical imaging modalities, such as CT and MRI, produce image data in slices, and these are conventionally viewed as an array of 2D grey-scale slice images. The advantages of comprehension in three dimensions have spurred the development of methods of 3D display of this data. 3D display of single, high contrast objects, such as bones imaged by CT, is beginning to be widely used in both diagnosis and surgical planning. Moreover, early systems for the display of lower contrast objects, such as soft tissue imaged by CT, and of the relationships of multiple objects, both anatomic and computed, show considerable promise for radiotherapy treatment planning, surgical planning, and diagnosis. This paper will survey approaches to effective 3D display to achieve these ends.

3D display can be roughly divided into two approaches. In the first, presently more common, approach a surface is determined from the original recorded image slices, and reflections of light impinging on this surface are simulated. We will call this approach *slice-based reflective display*. In the second approach the display is produced directly from the original recorded intensities, thought of as volume elements (*voxels*). Some combination of transparency, reflectivity, and (possibly colored) luminosity is attributed to each voxel, and the transmission through and reflection by these voxels from a chosen set of light sources is simulated. We will call this approach *voxel intensity based display*. We will focus on methods of slice-based reflective display because of their much more common use, but we will follow their description with a short summary of voxel intensity based methodology. We will also spend some attention on visualization methods that are applicable to both types of display.

Object Definition

Slice-based reflective methods require a preliminary stage of definition of the surfaces of the objects to be displayed, and voxel intensity based methods benefit from such object definition. The objects to be defined include anatomic objects, such as bones, internal organs, tumors, and skin, and treatment-related objects such as surgical incision planes, implants, stereotactic probes, and radiotherapy treatment beams and isodose surfaces. Unfortunately, methods that are satisfactory for the definition of anatomic objects from the original recorded slice images, whether by automatic or interactive means, are yet to be demonstrated, except for high contrast objects. For other anatomic objects the methods available are either far too prone to error or far too slow.

The available object definition approaches involve two steps: 1) the definition of either (i) 3-space points that are in the object (the *figure*) or (ii) points on the object surface, possibly with some information defining inter-point connections; and 2) the specification of a surface from these points. Both of these steps are somewhat problematical. Rather than list hopeful approaches for definition of anatomic objects, a most important area of research if 3D display is to become widespread, we will describe the only approach to the definition of surface points that is in common clinical use and then discuss a few surface specification methods that have proven useful.

Specification of figure or surface pixels

The only object definition techniques that seem satisfactorily easy to use are those involving a combination of intensity thresholds and user drawing on the original image slices. For high contrast objects it is frequently possible to define a range of recorded intensities such that on each slice, with few pixels excepted, the object is fully defined. That is, if figure pixels are being specified, the object forms a connected region of pixels with values inside the range, or if boundary pixels are being specified, a connected boundary between pixels with values inside the range and those with values outside the range forms the object boundary on the slice. The selection of the appropriate range is done interactively using a form of 2D (grey-scale or color) display of the slices in which the in-range pixels are differentiated from out-of-range pixels. For example, those pixels that are in the selected range may be especially intensified or colored (see Figure 1).

It is normally necessary to augment range-related specification with a means of editing the resulting set of figure or boundary pixels. For boundary drawing, specification of a piecewise linear section using a mouse or other pointing device to define successive boundary points is a common means. For editing the figure pixels, painting with such a device is a common means. Another useful editing tool is a program which chooses a smallest adequately fitting subset from an original set of boundary pixels, where the sequence of boundary pixels is taken to define a piecewise linear boundary. An effective algorithm of this type [1] operates by starting with a few extreme pixels from the original set and successively adding the pixel from the original set that is farthest from the boundary defined by the pixels that have been so far included. This addition of the "most needed" pixel continues until all remaining points in the orginal boundary set are within some predefined distance from the boundary so far defined or until some number of points have been chosen.

Surface specification

If the result of the initial step is a set of pixels making up the object in each slice, two approaches have had some success. In the first the pixels are formed into voxels by extending them in the inter-slice direction to abut those from the adjacent slices. Then from the union of the voxels from these slices one can compute the voxel faces such that the face-adjacent voxel is not in the set. Artzy, Frieder, and Herman [2] have reported an elegant algorithm to find this set of surface voxel faces, in such a way that faces that are edge-adjacent are connected in the resulting data structure.

In the second approach [3] a decision function, defined at any point in 3-space and giving a binary result of whether the point is inside or outside the object, is determined from the list of pixels in the two nearest slices. A figure with a smooth surface can thus be produced.

If the result of the initial step is a set of boundary contours on each slice, a skin must be placed over these contours to form the object surface. Two approaches can be distinguished. In the first [4] the boundary contour is taken to define a ribbon parallel to the inter-slice direction and of height equal to the distance between contours. A surface with steps between slices is thus defined. The second approach [5] defines a smoother surface that is piecewise polygonal (see Figure 2). In the Fuchs-Kedem-Uselton method the surface is computed as that of minimum area, by minimizing a path through a graph. This approach has difficulty with branching surfaces and those for which the contours on adjacent slices are non-overlapping. However, we have developed automatic means augmented by interactive editing for defining which contours are to be connected, where saddle points should lie, and what constraining polygon edges connecting contour points must be defined [6]. Given as input these constraining connections, the Fuchs algorithm routinely computes the surface tiling.

Rendering of Surfaces

Computing and displaying the intensity (and color) of each screen pixel so as to give the impression of a solid object is called *rendering*. While rendering single, simple objects so that they appear opaque can provide adequate comprehension, the appreciation of complex surfaces and especially of the 3D relationships among multiple anatomic objects or among anatomic objects and superimposed, computed treatment-related objects requires the rendering of objects as transparent (see Figures 3a and 3b).

Rendering begins with a specified viewpoint and involves determining which surface(s) are visible at each screen pixel, calculating for each such surface its own intensity at that pixel, and combining these intensities if more than one surface is simultaneously visible at that pixel (due to transparency). The intensity of a surface at a point is intended to convey its depth. This can be done by depth-shading, in which intensity decreases as depth from the viewer increases. However, the most successful display methods, even if they use this depth cue, also use a cue of the variation of reflected light with surface angle. This angle is represented by the vector normal to the surface at the display pixel in question.

In the common methods that represent the surface by a collection of planar polygons the normal direction, or at least the shading, must be interpolated across the polygons if the surface is not to appear piecewise planar (see [7] for a modern, general formulation). The simpler approach of interpolating the shading is shared by the widely used method of Gouraud [8,9] and that of Udupa [10]. In Gouraud shading linear interpolation of shading across each polygon is done, from shadings computed at the polygon's vertices. These vertex shadings, in turn, are computed by averaging the normals to the polygons impinging upon that vertex. In Udupa's method (see Figure 3b, for a result) the polygons are the surfaces of rectangular voxels, so they are limited to six orientations. The normal at any voxel face is interpolated from the orientation of that face and the four edge-adjacent faces.

The major method involving true interpolation of the normal is called *Phong shading* [9,11]. It has the advantage of allowing specular reflections to be included. This is especially helpful in visualizing the surfaces when multiple light sources are simulated and many surfaces, with some transparent, are to be simulataneously portrayed (see Figure 3a).

Visualizations and Viewpoints

Rendering requires the pre-specification of the viewpoint, but comprehension is frequently strongly dependent on the viewpoint, and often on providing images from multiple viewpoints. But the determination of the proper viewpoint(s) requires a visualization of the objects to be rendered, so a means of rough visualization with interactively modifiable viewpoint may be required. Vector (piecewise linear) display of the defining slice contours or the surface tiles may be provided for this purpose (see Figure 2).

Considerably increased comprehension of 3D relationships can be obtained if additional 3D cues besides obscuration and shading are provided. The most common of these require rendering from more than one viewpoint. Two such cues can be provided by pre-computation: stereoscopy and the kinetic depth effect. Another, interactive change of viewpoint, requires fast on-line computation.

To provide the stereoscopic effect, one renders the image from two viewpoints that differ by a few (~5) degrees horizontally, and one presents one of these to the left eye and the other to the right eye, with little enough delay between the presentations that they are perceived as simultaneous. Numerous means have been developed to do this approximately simultaneous presentation, but many have limitations by an expensive or burdensome apparatus for the viewer or limitations on the screen brightness, the number of displays simultaneously viewable, or the number of simultaneous viewers. One means has none of these limitations, the polarizing plate. This commercially available plate [12] is placed over the viewing screen, and it can switch in a few microseconds from passing light with one polarization to passing another. The viewer wears inexpensive, passive spectacles with each eye passing a different one of the polarizations. It is then only necessary to signal the plate to switch polarizations between successive presentations of the left eye's image and the right eye's image.

The kinetic depth effect is achieved when the displayed objects are made to appear to rotate in 3-space. It is obtained by computing the rendered image from viewpoints separated by a few degrees, and then displaying these rendered images in succession. While 360° rotation is preferred for giving a wide range of viewpoints, it requires the computation of tens of renderings. Rocking over a few degrees can be provided by as few as five renderings, and if the range of viewpoints is adequate, a helpful 3D cue is provided. However, this cue is weakened by the fact that as soon as the motion is stopped, to allow pointing or careful examination, the cue is lost.

Viewer-controlled dynamic specification of viewpoint has been shown to give much improved 3D comprehension over the stereopscopic or kinetic depth effects, in which the viewpoint is pre-specified. Furthermore, dynamic viewpoint specification can be combined with stereoscopy. The difficulty with providing such viewpoint specification has been that the very great computing power needed to render the image in a fraction of a second has not been available. However, such display engines are just now becoming available. One is Pixel-planes [13], a device in which a binary tree of 1-bit adders and a 1-bit adder at each pixel position in a 512 x 512 image buffer allow all pixels (x,y) to compute their own value of any linear function of x and y in 8 bits per microsecond. The greatly increased comprehension provided by dynamic viewpoint specification together with the abilities to dynamically vary other display properties such as the objects displayed, the objects' transparencies and colors, and the region of space viewed will probably make devices of such power popular as they become commercially available for less than $100,000.

Voxel Intensity Based Approaches

It has been long realized that slice-based reflective approaches had a significant weakness: 3D visualization depended on object defintion, viewpoint specification, and frequently contour connections, the optimal specification of all of which depended on 3D visualization. Numerous investigators therefore tried to develop methods of display that came directly from the voxel intensities. For example, Harris [14] combined the projection of voxel intensities with the kinetic depth effect, and Baxter [15] truly displayed intensities in 3-space using the varifocal mirror. But the projective method foundered on the inability to compute the projections fast enough and on the obscuration of objects of interest by other regions of 3-space in front of or behind them. The true 3D display methods also had these obscurational problems and were further weakened by the limited number of 3-space points that could be displayed. With the advent of much more powerful computing and display engines, such as the Pixar Image Computer [16] and Pixel-planes, the technical limitations have been much lessened, so encouraging pilot work on limiting the obscuration has begun.

Two important notions underly this work. First, for each object of interest the probability of each voxel being in that object can be computed from the voxel intensity and its relation to a few neighboring voxels, with these probabilities being accurate enough that a presentation based on them will allow at least rough visualization of the object. Second, an effective presentation can be based on simulating light transmission and associating color, luminosity, transparency, and reflectivity to each voxel based on its probabilities of being in each object or in the case of reflectivity, the directional derivatives of the probability. With appropriate choice of colors and transparencies that can be interactively varied, perhaps together with a means interactively to specify viewpoint and the part of space of interest (and perhaps improved 3-space object definition techniques), it appears that presentations competitive with the slice-based ones can be achieved (see Figure 4). With the powerful display engines that are reasonable to expect, however, these presentations could be produced on the spot from the original image data without the tedious slice-by-slice pre-specifications necessary with the slice-based reflective methods.

2D Adjunctive Display

No 3D display method can retain the ability to communicate the subtle grey-scale variations that 2D display can provide. At the same time, the display of an array of slices cannot provide the 3D comprehension achievable with 3D display. It is therefore useful to provide a tool for understanding grey scale spatial variation in a 3D context. Höhne [17] has demonstrated the importance of providing context for the grey-scale information by his display in which the parts of an object behind a clipping plane are presented by the reflective approaches discussed above but the clipped surface of the object shows grey-scale variations (see Figure 5). These grey-scale variations may be those of voxels on the clipped surface itself or the projection of recorded intensities within a fixed distance behind that surface. The newer display engines might be able to provide display of this form that is fast enough to keep up with interactive variation of this clipping plane.

Summary and Medical Use

Slice-based reflective 3D display has demonstrated its clinical usefulness for the display of single objects that appear at high contrast in the original slice images. While only a few strictly diagnostic applications have been provided by such display, it has provided important improvements for planning of bone surgery [18-20]. Improved contrast for other organs, e.g., from MRI, or improved object definition techniques would lead to usefulness in the planning of their surgery.

Improved 3D visualization of multiple objects provided by transparency augmented by multiple lights and Phong shading and additional depth cues, including interactively chosen viewpoints, can allow the visualization of many anatomic objects and thus improve surgical planning . These approaches can provide the appreciation of the relationship of treatment beams and isodose surfaces to patient anatomy [21-24] that is crucial for radiotherapy treatment planning. Pilot systems for these applications are already illustrating their worth.

Widespread application of 3D display for radiological diagnosis probably awaits the development of interactive display methods directly from voxel intensities. Therapy planning applications will also be improved by such development.

Acknowledgements

The preparation of this paper and the work on methods developed at the Univ. of North Carolina was done with the partial support of NIH grant # R01-CA39060.

We gratefully acknowledge our clinical co-investigators, Drs. Julian Rosenman, Jeff Creasy, and Jordan Renner, and our physics, programming, and engineering colleagues, Sandra Bloomberg, Edward Chaney, John Gauch, Amerie Helton, Mark Levoy, Peter Litwinowicz, Peter Mills, Charles Mosher, Kevin Novins, Richard Rubin, George Sherouse, Philip Stancil, and Lee Westover. We appreciate the photographs contributed by Karl-Heinz Höhne of the University of Hamburg and Mark Levoy of UNC. We thank Sharon Laney for help in preparing this manuscript and Bo Strain and Karen Curran for photography.

References

1. Ballard, D.H.; and Brown, C.M. (eds.) Computer Vision. Englewood Cliffs, New Jersey: Prentice-Hall, Inc. 1982 {Algorithm 8.1, 239}.
2. Artzy, E.; Frieder, G.; Herman, G.T.: The theory, design, implementation and evaluation of a three-dimensional surface detection algorithm. Comput. Graphics Image Process. 15 (1981) 1-24.
3. Dekel, D.: Fast, high quality, medical 3D software system. Medical Imaging, SPIE Proceedings, Newport Beach, California, February 1-6, 1987.
4. Heffernan, P.: A new method for shaded surface display of biological and medical images. IEEE Trans. Med. Imaging, MI-4(1) (1985) 26-38.
5. Fuchs, H.; Kedem, Z.M.; and Uselton, S.P.: Optimal surface reconstruction from planar contours. Communications of the ACM, 20(10) (1977) 693-702.
6. Pizer, S.M.; Fuchs, H.; Mosher, C.; Lifshitz, L.; Abram, G.D.; Ramanathan, S.; Whitney, B.T.; Rosenman, J.G.; Staab, E.V.; Chaney, E.L.; Sherouse, G.: 3D shaded graphics in radiotherapy and diagnostic imaging. Proc. Computer Graphics '86, National Computer Graphics Association, III (1986) 107-113.
7. Kajiya, J.T.: The rendering equation. Computer Graphics 20(4) (1986) 143-149 {SIGGRAPH '86}.
8. Gouraud, H.: Continuous shading of curved surfaces. IEEE Transactions on Computers, C-20(6) (1971) 623-628.
9. Foley, J.D. & van Dam, A. (eds.). Fundamentals of Interactive Computer Graphics. Reading, Mass: Addison-Wesley (1982). .
10. Chen, L.S.; Herman, G.T.; Reynolds, R.A.; Udupa, J.K.: Surface shading in the cuberille environment. IEEE Computer Graphics & Application (12) (1985) 33-43.

11. Bui-Tuong, P.: Illumination for computer-generated pictures. Communications of the ACM, 18(6) (1975) 311-317.
12. Bos, P.J.; Johnson, Jr., P.A.; K. Rickey Koehler/Beran, Tektronix, Inc.: A liquid-crystal optical-switching device (π cell). Proc. Soc. for Infor. Disp. (1983) 30.
13. Fuchs, H.; Goldfeather, J.; Hultquist, J.P.; Spach, S.; Austin, J.D.; Brooks, F.P.; Eyles, J.G.; and Poulton, J.: Fast spheres, shadows, textures, transparencies, and image enhancements in pixel-planes. Computer Graphics 19(3) (1985) 111-120, {SIGGRAPH '85, 29-37}.
14. Harris, L.D.; Robb, R.A.; Yuen, T.S.; Ritman, E.L.: Noninvasive numerical dissection and display of anatomic structure using computerized x-ray tomography. Medical Imaging 152, Proc. SPIE (1978).
15. Baxter, B.; Hitchner, L.E.; Anderson, R.E.: Application of a three-dimensional display in diagnostic imaging. J. Comp. Ass. Tomo. 6(5) (1982) 1000-1005.
16. Levinthal, A.; Porter, T.: Chap-A SIMD Graphics Processor. Computer Graphics 18(3) (1984) 77-82 {SIGGRAPH '84}.
17. Höhne, K-H.; et. al.: Three-dimensional investigation of tomographic volumes. Mathematics and Computer Science in Medical Imaging (Proc. NATO Advanced Study Institute, Il Ciocco, 1986) to appear in Springer-Verlag, 1987.
18. Hemmy, D.C.; David, D.J.; Herman, G.T.: Three-dimensional reconstruction of craniofacial deformity using computed tomography. Neurosurgery 13 (1983) 534-541.
19. Vannier, M.W.; Marsh, J.L., Warren, J.O.: Three dimensional CT reconstruction images for craniofacial surgical planning and evaluation. Radiology 150 (1984) 179-184.
20. Burk, D.L.Jr.; Mears, D.C.; Kennedy, W.H.; Cooperstein, L.A.; Herbert, D.L.: Three dimensional computed tomography of acetabular fractures. Radiology 155 (1985) 183-186.
21. Goitein, M.; Abrams, M.: Multidimensional treatment planning. I: Delineation of anatomy and II: with Rowell, D.; Pollari, H.; Wilaes. J.: Beam's eye-view, back projection, and projection through CT sections. Int. J. Radiation Biol. Phys. 9 (1983) 777-797.
22. Bauer-Kirpes, B.; Schlegel, W.; Boesecke, R.; Lorenz, W.J.: Display of organs and isodoses as shaded 3-D objects for 3-D therapy planning. Int. J. Radiation Oncology Biol. Phys. 13 (1987) 135-140.
23. McShan, D.L.; Silverman, A.; Lanza, D.M.; Reinstein, L.E.; Glicksman, A.S.: A computerized three-dimensional treatment planning system utilizing interactive colour graphics. Brit. J. Radiol. 52 (1979) 478-481.
24. Chaney, E.; Rosenman, J.; Sherouse, G.; Bourland, D.; Fuchs, H.; Pizer, S.; Staab, E.; Varia, M.; Mahaley, S.: Three dimensional displays of brain and prostate implants. Endocurie/Hyperthermia Oncology 2 (1986) 93-99.

1. **Threshold-Based Coloring in Interactive Object Definition of a Pelvic Bone on a CT Slice**

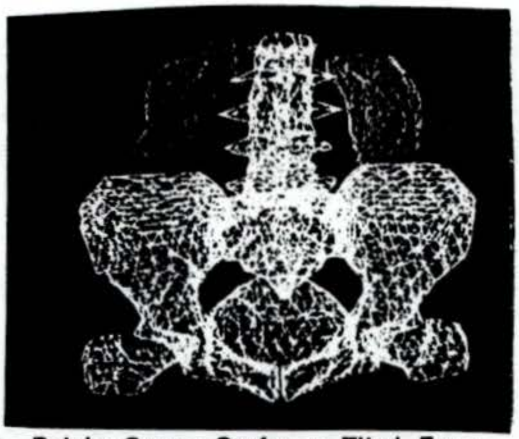

2. **Pelvic Organ Surfaces Tiled From Slice Contours Using Fuchs-Kedem-Uselton Method**

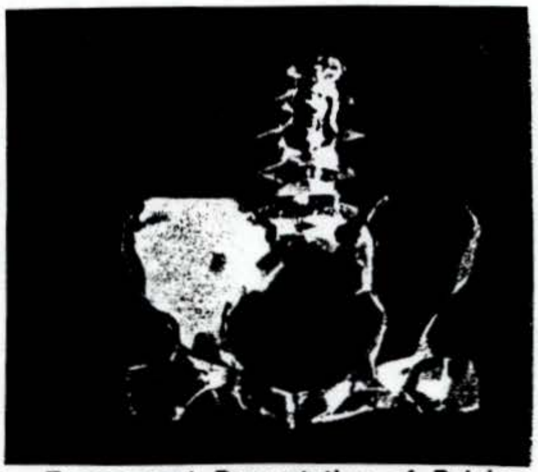

3a. **Transparent Presentation of Pelvic Organs Using Phong Shading and Multiple Lights**

3b. **Transparent Presentation of Tumor and Skull Using Interpolated Shading**

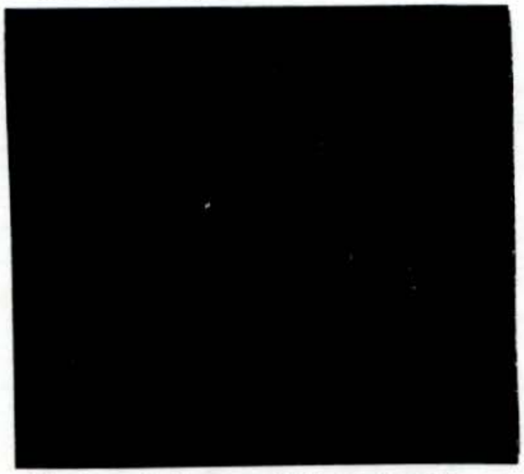

4. **Display of Face and Skull Based Directly on CT Voxel Intensity Gradients. Courtesy of M. Levoy**

5. **Grey-Scale Slice Superimposed on Reflective Surface Display of Head Courtesy of K. Höhne**