

Volume Rendering in Radiation Treatment Planning

*Marc Levoy^{1,2}, Henry Fuchs^{1,2}, Stephen M. Pizer^{1,2,3},
Julian Rosenman^{1,2}, Edward L. Chaney², George W. Sherouse²,
Victoria Interrante¹, Jeffrey Kiel⁴*

¹Computer Science Department
University of North Carolina
Chapel Hill, NC 27599

²Department of Radiation Oncology
North Carolina Memorial Hospital
Chapel Hill, NC 27599

³Department of Radiology
North Carolina Memorial Hospital
Chapel Hill, NC 27599

⁴School of Arts and Sciences
University of North Carolina
Chapel Hill, NC 27599

Abstract

Successful treatment planning in radiation therapy depends in part on understanding the spatial relationship between patient anatomy and the distribution of radiation dose. We present several visualizations based on volume rendering that offer potential solutions to this problem. The visualizations employ region boundary surfaces to display anatomy, polygonal meshes to display treatment beams, and isovalue contour surfaces to display dose. To improve perception of spatial relationships, we use metallic shading, surface and solid texturing, synthetic fog, shadows, and other artistic devices. Also outlined is a method based on 3D mip maps for efficiently generating perspective volume renderings and beam's-eye views. To evaluate the efficacy of these visualizations, we are building a radiotherapy planning system based on a Cray YMP and the Pixel-Planes 5 raster display engine. The system will allow interactive manipulation of beam geometry, dosimetry, shading, and viewing parameters, and will generate volume renderings of anatomy and dose in real time.

1. Introduction

Failure of radiation therapy to control the local-regional component of a malignancy can be due to misregistration of the radiation beams with the tumor or the use of radiation source arrangements yielding isodose surfaces that do not conform well to the target volume. Inaccurate targeting and poor isodose conformation may in turn be due to the physician's inability to appreciate the 3D shapes and spatial relationships of the tumor, target volume, normal anatomy, radiation beams, and dose distribution.

It has long been recognized that computer-generated 3D visualizations might be an effective means for presenting treatment planning data to the physician [20]. Although clinical trials comparing 3D planning methods with traditional 2D methods are

lacking, comparison on an anecdotal basis suggests that 3D methods often result in better plans. Accordingly, treatment planning systems are currently being developed at several institutions [5, 6, 2, 9, 16, 25, 18, 17]. We at the University of North Carolina have investigated both wire-frame and surface-based 3D visualizations [21] and have developed an interactive radiotherapy planner based on these visualizations [24].

Although wire-frame and surface-based visualizations are effective for displaying geometrically-defined objects, they have not proven entirely satisfactory for displaying anatomy or dose. This is in part because surface rendering requires that each anatomic feature be contoured, a time consuming and labor intensive process. The resulting polygonal mesh not only omits potentially useful information, but also creates the false illusion that features have well defined surfaces. For radiation dose, automatic contouring of isodose contour surfaces is possible, but large numbers of polygons must be used to produce a smooth mesh, and each dose level to be displayed must be contoured separately.

Volume rendering is a family of methods for visualizing sampled scalar or vector fields of three spatial dimensions without fitting geometric primitives to the data. A subset of these methods generates images by computing a color and a partial opacity for each voxel and then blending together contributions made by voxels projecting to the same pixel on the picture plane [11, 3, 23, 26]. Quantization and aliasing artifacts are reduced by avoiding thresholding during data classification and by carefully resampling the data during projection. The principal advantages of these techniques over wire-frame or surface-based visualizations are their superior image quality and the ability to generate images without explicitly defining surface geometry.

In this paper, we explore the application of volume rendering to radiation treatment planning. Our visualizations incorporate multiple datasets, mixtures of polygon and volume data, and combinations of 2D and 3D imagery. The visualizations also incorporate a number of artistic devices designed to

improve rendition of the data and facilitate understanding of the relationships between datasets. These devices are common in hand-drawn scientific and medical illustrations, but have not been applied to computer-generated medical imagery. Algorithms for rendering many of these devices already exist in the realistic image synthesis literature; others must be invented anew for volume rendering.

2. Display of anatomy, dose, and beams

The volume rendering method used in this paper is a hybrid ray tracer capable of handling both polygon and volume data [14]. It begins with a 3D array of voxel data. The array is shaded and classified to yield a color and an opacity for each voxel. Viewing rays are then traced into the array from an observer position. For each ray, samples are drawn along the ray, and a color and opacity is computed at each sample position by trilinearly interpolating from the colors and opacities of the nearest eight voxels. Independently, all intersections between the ray and polygons in the environment are computed and shaded, yielding a color and opacity for each point of intersection. The resampled volume colors and opacities are composited with each other and with the polygon colors and opacities in depth-sorted order to yield a color for the ray.

Depending on the method used to calculate voxel opacity, this rendering method can give volume data the appearance of opaque surfaces, semi-transparent surfaces, or semi-transparent volumes. In choosing among these, we have allowed ourselves to be guided by a few observations concerning human visual perception. Our world is dominated by surfaces. We are better at evaluating the shapes and spatial relationships of objects from reflections off their bounding surfaces than from transmission or scattering of light through their interiors. Most of the objects we encounter in daily life are also opaque. The difficulty of comprehending semi-transparent objects is evident to anyone who has visited the Steuben museum of glass sculpture in New York City.

In light of these observations, we have chosen to use opaque region boundary surfaces, i.e. surfaces bounding tissues of constant density, to display anatomy. For similar reasons, we have chosen to use opaque or semi-transparent isovalue contour surfaces, i.e. surfaces defined by points of tissue receiving the same energy, to display dose. Finally, we have chosen to use semi-transparent polygons to display treatment beams and other geometrically-defined objects. Using these visualizations, the spatial relationship between data elements is conveyed by the interpenetration of multiple surfaces. The opacity calculations required to volume render region boundary surfaces and isovalue contour surfaces are described in [11]. Suitable shading models for polygons are surveyed in [7].

3. Artistic devices

Metallic surface shading. The intensity of light reflected from a surface depends on the orientation of the surface with respect to the observer and light source and on the properties of the surface. Shiny surfaces exhibit a greater change in reflected light intensity for small variations in surface orientation than dull surfaces. Polished metals constitute a large and familiar class of shiny surfaces. For these reasons, many of our images simulate the appearance of polished metal. The clarity evident in the convolutions of the cortical surfaces in figure 7 demonstrates the utility of this effect. We approximate metallic shading by using a large exponent in the calculation of specular reflection in our

Phong shading model. More accurate shading models for metallic surfaces exist [1], but the improvement in comprehension of surface shape provided by the more expensive models would probably be small.

Embedded backdrops. Chemists working with brass models of complicated molecules frequently slip pieces of white paper behind groups of atoms they are studying. In so doing, they reduce the visual complexity of the image they see by limiting the depth complexity of the scene they are looking at. By analogy, the embedding of polygons in volume data improves comprehension of the latter. Figure 1 shows a $256 \times 256 \times 113$ voxel CT study of a human head. Five colored slabs have been embedded in the volume data using 3D scan-conversion with anti-aliasing [12]. The presence of backdrops of known shape and distinct colors enhances appreciation of the orbital bones.

Shadows. Illuminated objects in the real world cast shadows. If the shadows fall onto a surface of known geometry, such as a piece of paper inserted in a molecular model, the shadows help us determine the shape of the illuminated objects. Casting shadows of volumetrically-defined objects onto backdrops embedded in the volume data has a similar effect. The scene shown in figure 1 contains two light sources: a high-intensity light shining up from below and a low-intensity fill light shining over the observer's shoulder. To insure that the shapes of shadowed objects are not completely obscured, shadows were only computed for the high-intensity light. While this is not strictly correct, our goal is enhanced insight, not photorealism. Initial light strengths were assigned from a texture containing a filtered rectangular grid. The effect is to project this texture through the data and onto all illuminated surfaces. Shadows cast by the orbital bones on the transverse backdrop help us to interpret their shapes. A description of the shadow casting algorithm is contained in [12].

Cutting planes. Embedding backdrops in a scene reduces its depth complexity by obscuring the scene's backmost elements. Cutting planes reduce depth complexity by removing the frontmost elements. They also present a planar cross section of the scene in its 3D context. Most volume rendering systems provide some ability to display and manipulate cutting planes. Figure 2 shows four possible interpretations of a single cutting plane lying parallel to the image plane and embedded partway back through an $256 \times 256 \times 109$ voxel MR scan of a live human subject.

In the upper-left image, the opacity of all voxels lying in front of the cutting plane has been set to zero following calculation of colors and opacities. No shading has been applied to the cutting plane itself. Since voxel opacity is proportional to local gradient magnitude, organ boundaries are opaque and organ interiors are transparent. In the presence of a cutting plane, anatomic structures appear hollow, yielding a confusing image.

In the upper-right image, a similar treatment of the volume data has been used, but this time the raw voxel values have been displayed on the cutting plane using texture mapping. To prevent the applied texture from obscuring the volume rendering behind it, the texture was made transparent in voxels representing air. To help distinguish between the texture and the volume rendering, they are displayed in different colors. This interpretation of cutting planes is very common and is incorporated into many volume rendering systems [8, 10, 19].

In the lower-left image, no shading has been applied to the cutting plane, but the scalar value of all voxels lying in front of the cutting plane has been set to zero prior to calculation of colors and opacities. The resulting high gradient perpendicular to the cutting plane gives rise to additional surfaces during

shading, capping off anatomic structures and yielding a more natural image than the upper-left. Unfortunately, natural variations in tissue composition perturb the gradient vectors of voxels near the cutting plane, giving capping surfaces an uneven appearance.

In the lower-right image, a similar treatment is used except that voxels having values higher than a specified ceiling have been replaced with the ceiling prior to calculation of colors and opacities. This eliminates local perturbations of gradient vectors near the cutting plane, producing flat-looking capping surfaces. The contrast between curved anatomic structures and flat capping surfaces helps differentiate between them. The notion of remapping data values before shading to obtain a visual effect has also been reported by Drebin et al. [3].

In all of these methods, care must be taken to avoid aliasing artifacts. All but the upper-right image were generated by modifying the voxel array prior to ray tracing. To avoid artifacts, the transformation of voxel values was feathered out in a direction perpendicular to the cutting plane, producing a blurry edge in the voxel array. The cutting plane in the upper-right image was rendered during ray tracing. To avoid artifacts in this case, voxel values were carefully resampled during mapping onto the cutting plane [12], and the opacity of the applied texture was changed gradually from complete transparency to complete attenuation as voxels changed from air to tissue. Special treatment is also required to insure an artifact-free seam between the texture and volume rendered surfaces lying behind it [14].

Surface and solid texture. A single isovalue contour level usually presents only a fraction of the useful information in a volume dataset. Radiation oncologists, for example, need to visualize anatomic structures receiving dose falling within a specified range, contour levels from volume data is not easy. Figure 3 shows two isodensity surfaces from an electron density map of *Staphylococcus Aureus* ribonuclease. The shape of the semi-transparent outer surface and the spacing between the two surfaces is not obvious from this image. In figure 4, the opacity of the outer isovalue surface alternates between complete transmission and complete attenuation based on the value of a solid texture containing a honeycomb of small rectangular cells. To avoid aliasing artifacts, the solid texture is blurred slightly before use. In addition, all surface normals have been randomly perturbed before shading, giving the inner surface the appearance of hammered copper. These surface and solid textures improve our ability to distinguish the shape of the surfaces and their relationship to each other. This improvement is particularly marked in the two doughnut-shaped benzene rings at lower-left and upper-right.

Highlighted curves of intersection. A cartoonist drawing a ship in the water is certain to delineate the curve of intersection between the ship's hull and the waves. For two partially correlated volume datasets, curves of intersection between surfaces in the two datasets are equally revealing. The images in figure 5 show three opaque gray slabs scan-converted into a voxel array and two semi-transparent pink isovalue contour surfaces of a radially symmetric 3D field. In the left image, no special action has been taken where the slabs meet the isovalue surfaces. In the right image, the curves of intersection between slabs and isovalue surfaces have been highlighted in bright red. The spherical shape of the isovalue surfaces is more evident in the right image, as is their relationship with the slabs.

Synthetic fog. Given a 3D region enclosed by a semi-transparent surface and an object floating free in the region, it is difficult to determine the position of the object within the region. In the real world, fog provides a useful cue in such situations.

Figure 6 shows a $256 \times 256 \times 76$ voxel CT study of a male pelvis with a Foley catheter placed in the bladder. The dose distribution from four beams aimed at the prostate has been calculated using a generalized Clarkson scatter summation technique. A semi-transparent isodose contour surface has been displayed in red, and the region enclosed by the surface has been filled with a semi-transparent red fog. To facilitate study of the dose being delivered to the bladder, the right half of the dose distribution has been removed before shading.

4. Beam's-eye views

Clinical experience using our current radiotherapy planner has demonstrated the importance of providing users with beam's-eye views of patient anatomy [22]. These images allow ready identification of anatomic structures lying in the path of a planned treatment beam.

Since radiation beams diverge, beam's-eye views must be generated using a perspective projection. In volume rendering algorithms based on ray tracing, divergence causes the spacing of rays and hence between samples along rays to rise with increasing distance from the observer position. To avoid aliasing artifacts, care must be taken to avoid undersampling the data. We have developed a computationally efficient resampling algorithm based on 3D mip maps - the extension to three dimensions of the 2D mip map texture resampling method reported by Williams [27]. Our algorithm is described in a recently published paper on gaze-directed volume rendering [15].

Figure 7 shows a $256 \times 256 \times 109$ voxel MR scan of a live human subject. A polygonally defined target volume (in purple) and three treatment beams (in red) have been added using our hybrid ray tracer. A portion of the volume has been cut away and the raw MR values have been mapped onto the cutting planes. As in figure 2, the use of different colors for the 2D texture and 3D volume rendering facilitates distinguishing between them. Note that this is a sample visualization, not a clinical study; the MR data and treatment plan were taken from two different patients.

Figure 8 shows a right lateral beam's-eye view of the same subject. The placement of the beam with respect to the target volume and normal anatomy can be readily appreciated in this perspective view. To reduce visual confusion, polygons defining the opposing beam have been omitted.

5. Real-time interaction

Most researchers agree that routine use of volume rendering in the medical disciplines will be limited until images can be generated in real time or near-real time. A physician must be able to rotate a dataset, dynamically vary shading and classification parameters, and interactively move backdrops and cutting planes through the data. We are currently building an interactive radiation treatment planning system that provides all of these capabilities. It will also allow physicians to alter beam geometry and dosimetry and then view the resulting dose distribution superimposed over patient anatomy in real time. The system consists of a 4-processor Cray YMP, the Pixel-Planes 5 massively parallel raster display engine [4], and a variety of workstations and input devices.

The Cray YMP consists of 4 processors with a combined peak vector performance of 1200 MFLOPS and a large high-speed main memory (128 Mbytes). These characteristics make it ideal for radiation dose calculations, which require random access to 3D CT data. Pixel-Planes 5 consists of 16 40-MFLOP

graphics processors, 1/4 million pixel processors having a combined peak performance of 436,000 MIPS (8-bit integer adds), a 1024 x 1280 pixel color frame buffer, and a 640 Mbyte/sec ring network. The implementation on this machine of a near-real-time ray tracer for volume data has already been described [13]. The shading calculations for all voxels are performed in the pixel processors, and the ray tracing required to generate an image is divided among the graphics processors.

Using this strategy, we expect to obtain the fastest volume rendering system in the world by more than an order of magnitude as well as a radiation treatment planning system of unprecedented power and flexibility. The user will be able to interactively manipulate the treatment plan or viewing parameters and will receive feedback in the form of 512 x 512 pixel color images volume rendered from 128 x 128 x 128 voxel arrays at up to 8 frames per second. During periods of user inactivity, the resolution of the input arrays will refine in less than a second to 256 x 256 x 128 voxels.

6. Conclusions

We have presented a number of visualizations employing volume rendering that have potential application to radiation treatment planning. To improve the usefulness of our imagery, we have borrowed artistic devices from hand-drawn scientific and medical illustrations, and we have borrowed rendering techniques from the realistic image synthesis literature. We are currently implementing our system on a hardware platform of sufficient power to provide interactive control over the treatment plan and real-time generation of the resulting visualizations.

We wish to emphasize that the assertions made in this paper regarding the merit of specific visualizations are based on aesthetic intuition extrapolated from psychophysical principles. Our volume rendering-based radiation treatment planning system is not yet ready for clinical use. Hence, we have not yet performed controlled observer experiments to test our assertions. Even given such a system, the parameter space of possible volume rendered imagery is enormous, and we cannot expect to systematically evaluate more than a handful of visualizations.

Acknowledgements

The CT scans were provided by North Carolina Memorial Hospital, the MR study was provided by Siemens AG, and the electron density map was obtained from Dr. Chris Hill of the University of York. This work was supported by NCI grant P01-CA47982.

References

[1] Cook, R.L. and Torrance, K.E., "A Reflectance Model for Computer Graphics," *Computer Graphics*, Vol. 15, No. 3, August, 1981, pp. 307-315.

[2] Dahlin, H., Lamm, I.L., Landberg, T., Levernes, S., Ulso, N., "User Requirements on CT-Based Computed Dose Planning Systems in Radiation Therapy: Presentation of 'check lists,'" *Comp. Programs Biomed.*, Vol. 16, 1983, pp. 131-138.

[3] Drebin, R.A., Carpenter, L., and Hanrahan, P., "Volume Rendering," *Computer Graphics*, Vol. 22, No. 4, August, 1988, pp. 65-74.

[4] Fuchs, H., Poulton, J., Eyles, J., Greer, T., Goldfeather, J., Ellsworth, D., Molnar, S., Turk, G., Tebbs, B., and Israel, L., "A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," *Computer Graphics*, Vol. 23, No. 3, July, 1989, pp. 79-88.

[5] Goitein, M., Abrams, M., "Multi-Dimensional Treatment Planning: I. Delineation of Anatomy," *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 9, 1983, pp. 777-787.

[6] Goitein, M., Abrams, M., Rowell, D., Pollari, H., Wiles, J., "Multi-Dimensional Treatment Planning: II. Beam's-Eye View, Back Projection, and Projection Through CT Sections," *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 9, 1983, pp. 789-797.

[7] Hall, R., *Illumination and Color in Computer Generated Imagery*, Springer-Verlag, 1989.

[8] Hoehne, K.-H., Bomans, M., Tiede, U., and Riemer, M., "Display of Multiple 3D-Objects Using the Generalized Voxel-Model," *Proc. SPIE*, Vol. 914, 1988, pp. 850-854.

[9] Houliard, J.P., Dutreix, A., "3D Display of Radiotherapy Treatment Plans," *Proc. 8th ICCR*, 1984.

[10] Johnson, E.R. and Mosher, C.E., Jr., "Integration of Volume Rendering and Geometric Graphics," *State of the Art in Data Visualization*, Siggraph '89 tutorial, July, 1989.

[11] Levoy, M., "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*, Vol. 8, No. 3, May, 1988, pp. 29-37.

[12] Levoy, M., "Display of Surfaces from Volume Data," Ph.D. dissertation, University of North Carolina at Chapel Hill, May, 1989.

[13] Levoy, M., "Design for a Real-Time High-Quality Volume Rendering Workstation," *Proc. Chapel Hill Workshop on Volume Visualization*, ed. C. Upson, University of North Carolina, 1989, pp. 85-92.

[14] Levoy, M., "A Hybrid Ray Tracer for Rendering Polygon and Volume Data," *IEEE Computer Graphics and Applications*, Vol. 10, No. 2, March, 1990.

[15] Levoy, M., "Gaze-Directed Volume Rendering," *Proc. 1990 Utah Symposium on Interactive 3D Graphics.* In press.

[16] McShan, D.L., Fraas, B.A., Lichter, A.S., "Three Dimensional Portal Design and Verification" (abstract), *Med. Phys.*, Vol. 13, 1986, pp. 575.

- [17] Mohan, R., Barest, G., Brewster, L., Chui, C., Kutcher, G., Laughlin, J., Fuks, Z., "A Comprehensive Three-Dimensional Radiation Treatment Planning System," *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 15, 1987, pp. 481-495.
- [18] Nagata, Y., Nishidai, T., Abe, M., Takahashi, M., Yukawa, Y., Nohara, H., Yamaoka N, Saida, T., Ishihara, H., Kubo, Y., Ohta, H., Inamura, K., "CT Simulator: A New Treatment Planning System for Radiotherapy" (abstract), *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 13 (suppl.), 1987, pp. 176.
- [19] Robb, R.A. and Barillot, C., "Interactive Display and Analysis of 3-D Medical Images," *IEEE Transactions on Medical Imaging*, Vol. 8, No. 3, September, 1989, pp. 217-226.
- [20] Reinstein, L.E., McShan, D., Webber, B.M., Glicksman, A.S., "A Computer Assisted Three-Dimensional Treatment Planning System," *Radiology*, Vol. 127, 1978, pp. 259-264.
- [21] Rosenman, J., Sherouse, G.W., Fuchs, H., Pizer, S.M., Skinner, A.L., Mosher, C., Novins, K., Tepper, J.E., "Three-Dimensional Display Techniques in Radiation Therapy Treatment Planning," *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 16, 1989, pp. 263-269.
- [22] Rosenman, J.G., Sherouse, G.W., Chaney, E.L., Tepper, J.T., "Virtual Simulation: Initial Clinical Results," *Int. J. Radiat. Oncol. Biol. Phys.* Submitted for publication.
- [23] Sabella, P., "A Rendering Algorithm for Visualizing 3D Scalar Fields," *Computer Graphics*, Vol. 22, No. 4, August 1988, pp. 51-58.
- [24] Sherouse, G.W., Mosher, C.E., Novins, K., Rosenman, J., Chaney, E.L., "Virtual simulation: Concept and implementation," *The Use of Computers in Radiation Therapy, Proceedings of the Ninth International Conference*, 1987, pp. 433-436.
- [25] Smith, R.M., Galvin, J.M., Harris, R. "Correlation of CT Information with Simulation and Portal Radiographs" (abstract), *Int. J. Radiat. Oncol. Biol. Phys.*, Vol. 13 (suppl.), 1987, pp. 176.
- [26] Upson, C. and Keeler, M., "VBUFFER: Visible Volume Rendering," *Computer Graphics*, Vol. 22, No. 4, August 1988, pp. 59-64.
- [27] Williams, L., "Pyramidal Parametrics," *Computer Graphics*, Vol. 17, No. 3, July, 1983, pp. 1-11.

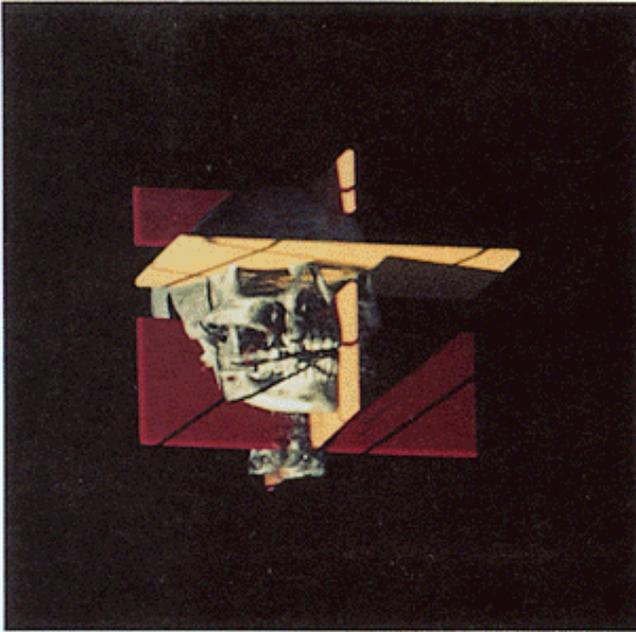


Figure 1: Volume rendering of $256 \times 256 \times 113$ voxel CT study of human head. Five colored slabs have been embedded in the volume data using 3D scan-conversion with anti-aliasing. Shadows have been cast from one of the two light sources in the scene using a two-pass algorithm. Initial light strengths were assigned from a texture containing a filtered rectangular grid, causing the texture to be projected onto all illuminated surfaces.



Figure 2: Volume renderings of $256 \times 256 \times 109$ voxel MR scan of a live human subject. A single cutting plane lying parallel to the image plane has been embedded in the volume data.
Upper-left: Frontmost voxels removed after shading.
Upper-right: Same. Raw voxels mapped onto cutting plane.
Lower-left: Frontmost voxels removed before shading.
Lower-right: Same. Voxel values remapped before shading.



Figure 3: Volume rendering of two isodensity contour surfaces from an electron density map of *Staphylococcus Aureus* ribonuclease. The inner surface is opaque and the outer surface is semi-transparent. Surface normals have been randomly perturbed before shading to provide surface texture.

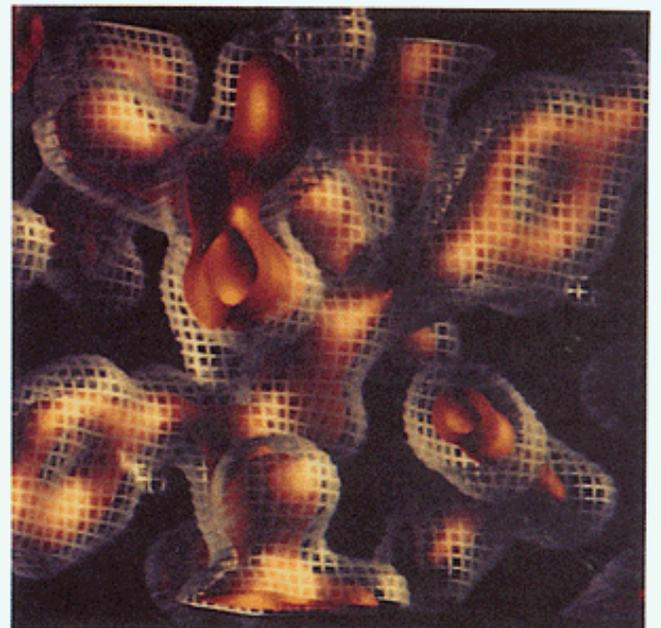


Figure 4: Same dataset. The opacity of the outer isovalue surface alternates between complete transmission and complete attenuation based on the value of a solid texture containing a filtered honeycomb of small rectangular cells.

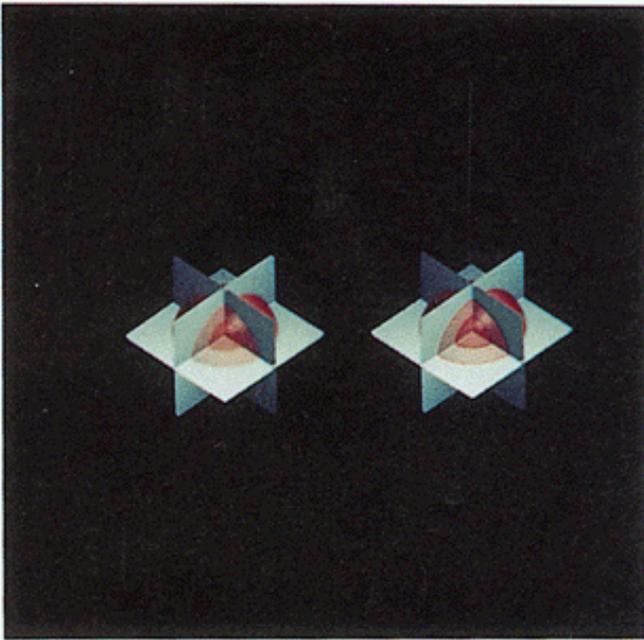


Figure 5: Volume renderings of opaque gray slabs scan-converted into a voxel array and semi-transparent pink isovalue contour surfaces of a radially symmetric 3D field.

Left: No special action at intersections.

Right: Intersections highlighted in bright red.



Figure 6: Volume rendering of $256 \times 256 \times 76$ voxel CT study of male pelvis with a Foley catheter placed in the bladder. An isodose contour surface has been displayed in red, and the region enclosed by the surface has been filled with semi-transparent fog.

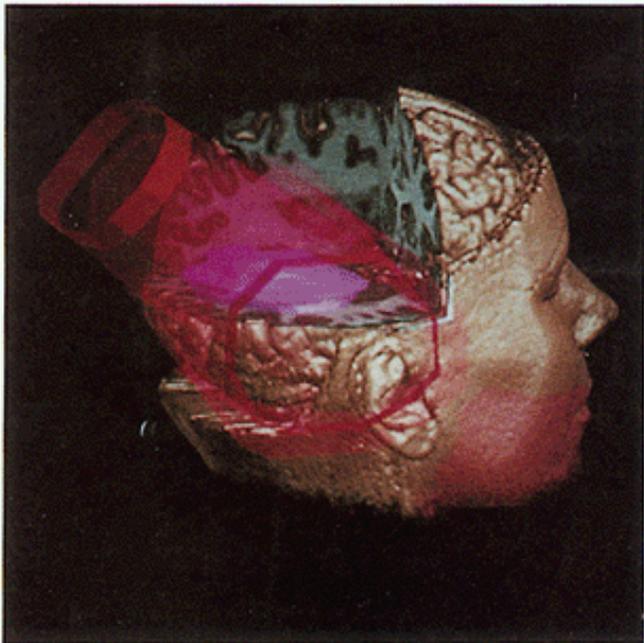


Figure 7: Volume rendering of $256 \times 256 \times 109$ voxel MR scan of a live human subject. A polygonally defined target volume (in purple) and three treatment beams (in red) have been added using a hybrid ray tracer. A portion of the volume has been cut away and the raw MR values have been mapped onto the cutting planes.



Figure 8: Right lateral beam's-eye view of the same subject. Note use of perspective projection. To reduce visual confusion, polygons defining the opposing beam have been removed from this view. These are sample visualizations, not clinical studies; the MR data and treatment plan were taken from two different patients.