Volume Rendering for Display of Multiple Organs, Treatment Objects, and Image Intensities

Stephen M. Pizer, Marc Levoy, Henry Fuchs, Julian G. Rosenman
Medical Image Display Research Group, University of North Carolina, Chapel Hill, NC 27599-3175, USA

ABSTRACT

Volume rendering not only produces high quality 3D images of single organs directly from the intensity data in CT, MR, SPECT, PET, and ultrasound images, but it also can be used to show the relationship among multiple anatomic, treatment, and image objects. In this paper we will explain methods, show results, and discuss the effectiveness of

1. simultaneously volume rendering multiple organs from a single image data set, using transparency and color;
2. rendering polygonally defined treatment objects such as prostheses and radiation treatment beams with volume rendered anatomy.

We will also discuss ongoing work in

1. rendering multiple image sets, e.g., radiation dose distributions and medical image data, into a single image via volume rendering, e.g., to show the relation between isodose surfaces and anatomic objects;
2. texture mapping grey scale slices onto clipping planes on volume rendered anatomy, to show the relationship between physiological and anatomic data, or to show the subtle anatomic intensity variations in its 3D context.

1. INTRODUCTION

While display of single objects, such as a bone or soft tissue organ, has been the most common objective of 3D medical image display, clinical practice frequently requires the portrayal of the relationship among anatomic objects, the relationship among anatomic objects and treatment objects, such as prostheses, incision planes, radiation treatment beams, and isodose surfaces, and the relationship between physiological function and anatomical location. Moreover, since 3D display normally portrays surfaces, it does not directly portray subtle or continuous variations in the intensities in the original 3D data set, e.g., from CT, or MRI (we will call these the *original image intensities*). Recovering this loss by providing comprehension of the rélationship of original image intensities and anatomic surfaces has been shown to be a clinically useful objective [Höhne, 1987].

Volume rendering [Levoy, 1988a; Drebin, 1988], which is based on surface likelihoods and involves no intermediate geometric representation such as tiles, is being shown to produce images of greater detail and accuracy than surface rendering techniques, which are based on explicit surface location and surface representation by tiles. At the same time the fact that, unlike surface rendering, volume rendering produces the image directly makes the whole process of starting from a 3D set of medical image data and producing a 3D image of a surface faster than with surface rendering, except for objects appearing in the original image data with such high contrast, sharp boundaries that fast, automatic determination of the object surface is possible. In fact, volume rendering has the potential of exploratory variation of viewing and surface classification parameters, with the image being produced in a fraction of a second on parallel graphics computers [Levoy, 1989].

In addition to its strengths for direct, high quality display volume rendering allows the presentation of multiple objects and of image intensity information with 3D objects. Our methods toward this objective, while still in formation, and some results to date are the subject of this paper.

2. METHODS

Volume rendering is based in principle on calculating at every voxel an opacity and a shade and then compositing these opacities and shades for each pixel by projection along a ray from the viewpoint through the pixel into the 3D opacities and shades, using the relation

$$C_{out}(\mathbf{u};\mathbf{U}) = (1 - \alpha(\mathbf{U}))\, C_{in}(\mathbf{u};\mathbf{U}) + \alpha(\mathbf{U})\, C(\mathbf{U}), \qquad (1)$$

where $C_{out}(\mathbf{u};\mathbf{U})$ is the displayable intensity coming out of location **U** on ray u toward the viewer, $\alpha(\mathbf{U})$ is the opacity at the ray location **U** (interpolated from nearby voxel opacities), $C(\mathbf{U})$ is the shade at that ray location (likewise interpolated), and $C_{in}(\mathbf{u};\mathbf{U})$ is the displayable intensity coming into the ray location from behind. The combination of multiple objects and or grey scale slices is based on compositing either from multiple sets of opacities and shades or from a single set of opacities and shades arranged to account for many objects. For example, one set of opacities and shades might come from a surface classification for one tissue type in an anatomic array of original intensities, another set from a different surface classification for a different tissue type in that anatomic array, a third set from a surface classification for an entirely different array of original intensities, a fourth set from a surface represented by geometric primitives and then transformed into opacities and shades, and a fifth set from intensities in either of the two original intensity arrays or yet another array.

In both our work and that described in [Drebin, 1988] the compositing involves choosing display parameters such as color, transparency, and surface specularity for each contributing data set and then, primary color component by component, combining all the data sets according to equation 1. That is, every ray location at which compositing occurs must undergo opacification by each of the opacities there from the various contributing data sets and inclusion of all of the shades there from each of these data sets. An alternative is the combining, for every voxel, of the opacities and shades there of the various data sets into a single opacity and shade there, followed by a single compositing. The combination of opacities and shades is done via equation 1, so the result varies with the order of application of the opacities and shades contributing in the compositing at a ray location or voxel.

Image quality is strongly affected by the care with which antialiasing is done, not only in rendering objects based on original intensities. but especially when objects based on geometric representations are transformed into opacities and shades.

3. APPLICATIONS AND RESULTS

Multiple anatomic objects from a single 3D intensity array can be shown effectively either with a single classification or with multiple classifications [Drebin, 1988]. In the first case, one uses shade functions that vary, e.g., in color, with some image property such as original image intensity, and one designs an opacity function of local intensities to produce high opacities for surfaces of each object, e.g., in each of a number of intensity ranges. In the second case separate classifications and shadings, and thus opacities and shades, are computed for each of the objects, and one composites the resulting opacity and shade distributions using different global opacity multipliers and colors for each component. An example is shown in figure 1, in which on the color CRT display, data from a single 256 x 256 x 113 CT data array is used to portray the flesh in semitransparent red and the bone in opaque yellow using a single classification. The approach appears to be equivalent to the approach based on computing tissue mixture fractions used in the Pixar [Drebin, 1988]. Note that the approach is applicable not only to multiple classifications of a single intensity data set but also to classifications of a number of adequately registered different intensity data sets, e.g., from a tumor shown by PET and a brain surface shown by MRI [Levin, 1989].

Figure 1 also illustrates the effective presentation of the relationship of anatomy to user-defined geometrically represented objects. Here the treatment beam (in blue on the CRT color display) is shown semitransparently intersecting the conically shaped treatment region (in opaque orange on the CRT color display) and avoiding the radiosensitive eye. The treatment beam was defined by the user as a single cross-section in the collimator plane and extended into geometric tiles. Rays are cast through the tiles, using adaptive supersampling to provide adequate antialiasing, and the result is composited with the anatomic and treatment region opacities and shades [Levoy, 1988b]. In the future it can be expected that manipulation of the beam position and cross-section will be done interactively with immediate feedback from the volume rendered image.

Figure 1 also illustrates the presentation of regions defined by contour drawing on the original grey scale image slices, a common basis for surface rendering techniques. The treatment region, which depends on clinical judgment and therefore cannot be portrayed directly from the CT image data, is based on contours drawn on a number of slices using our X-based contour specification program [Mills, 1989]. Minimum surface tiling of these contours [Fuchs,1977] was followed by the calculation of a supersampled set of opacities and shades and compositing with the other sets of opacities and shades for the anatomy and beams.

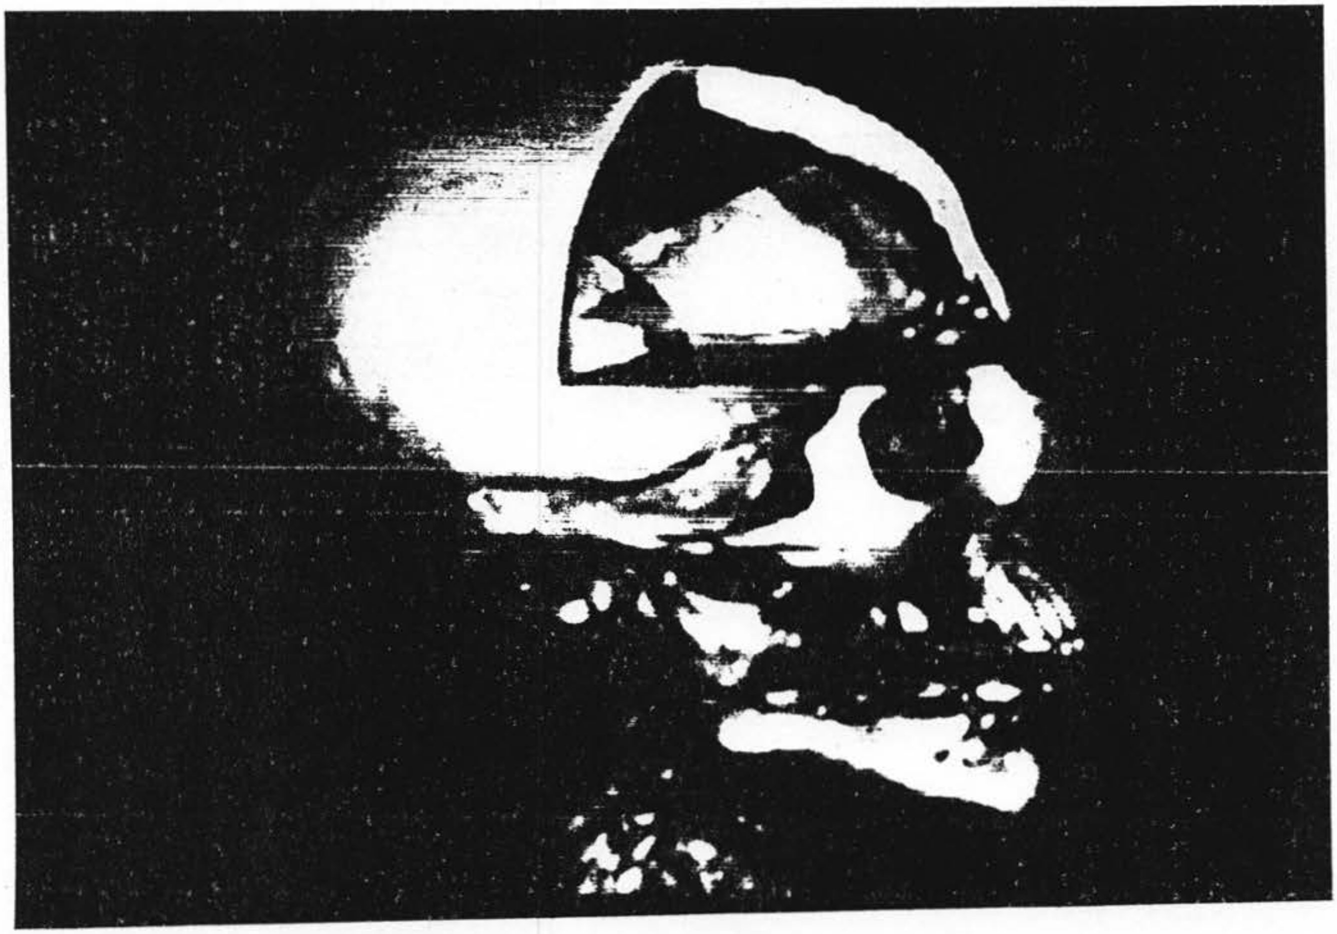

Figure 1. A black and white photo of a color rendering of flesh and bone of a head from CT data together with a radiation treatment beam and a treatment region from tiled representations.

Finally, figure 1 also illustrates the exposure of the treatment region by the removal of a exclusionary region of interest (ROI) defined as a simple geometric object. A similar technique, except with the exclusionary ROI specified so as to exclude a particular imaged anatomic feature, the skin, is illustrated in figure 2 for the presentation of the cortical surface of the brain from MRI. Here the bright region of the scalp could not be removed with the use of an adequately antialiased opacity function based on intensity values or intensity gradient magnitudes because the skin-air boundary crosses the same intensity levels as the brain-bone boundary. The skin region was therefore surrounded by a drawn contour, slice by slice, but note that this contour could be roughly drawn, only needing to fall within the black region between the scalp and the brain. Recent results of the use of regions defined via the Intensity Axis of Symmetry description of image intensity viewed as an intensity surface over 2-space [Gauch, 1988] suggest that one or two pointings per slice would suffice for the selection of the scalp exclusionary ROI. We are presently working toward extending this method to 3D, to the end that a total of only one or two mouse points would be necessary to define the whole 3D exclusionary ROI.

Geometric objects can be used not only for presentation of objects but also to aid visualization. Figure 3 shows the inclusion of background planes in bones volume rendered from CT data. We have noted that when a pair of such planes, one sagittal and one coronal, are opaquely or semitransparently presented in red behind the orbit, which is presented in grey, and the image is rerendered from a number of viewpoints rotationally separated by 3 degrees, and the resulting sequence is displayed as a cine loop, the appreciation of the structure of the thin bones in the orbit is much improved as compared to a similar presentation without the interposed background planes. Because the bones at the back of the orbit are thin, they appear in volume rendering with only moderate opacity, and the simple background of the planes allows better appreciation of their structure than an unknown and complicated background of the back of the skull. We suggest that interactive location of such background planes with immediate display providing feedback will be a useful approach.

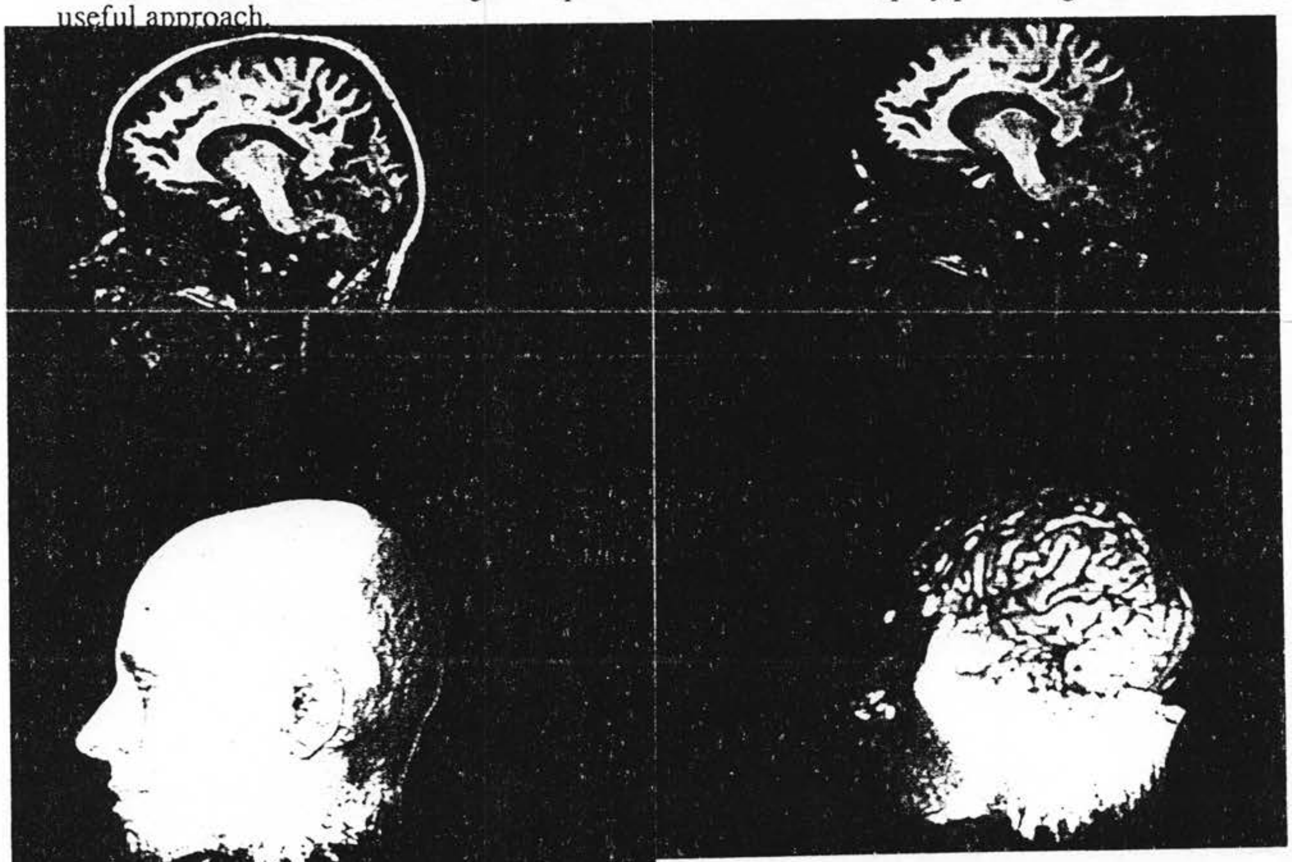

Figure 2. Volume rendering of a head surface (lower left) from a 3D MRI data set (a sample slice is shown on upper left). Editing of the slices to expose the cortex (upper right) produces a volume rendering of the brain (lower right).

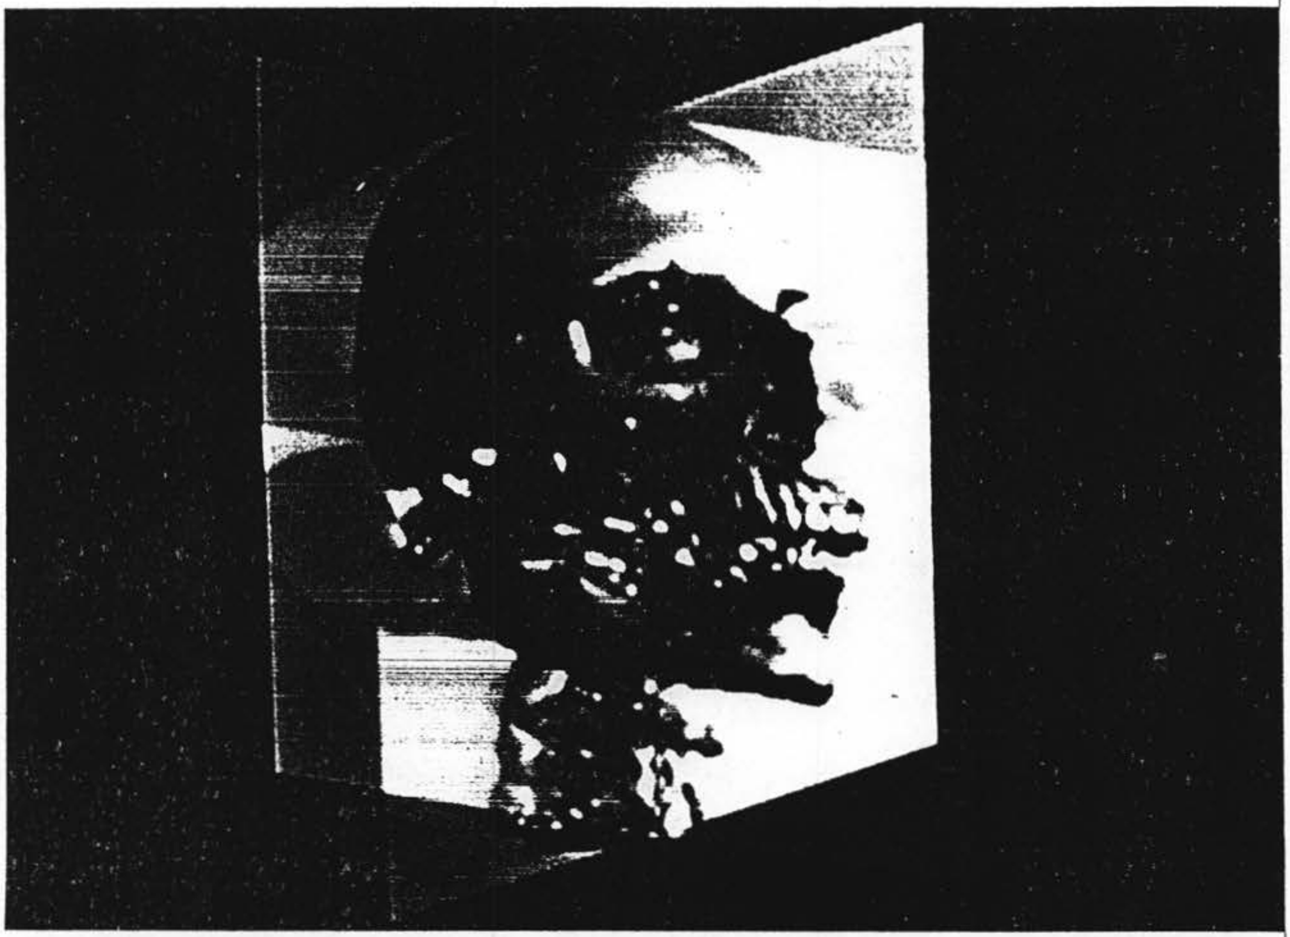

Figure 3. The use of geometrically defined planes to provide viewing background for enhancement of the rendition of moderately opacified regions. The image is based on CT of the head.

Work in progress suggests that compositing methods described above will allow the superimposition of grey level representations a slice of original or other image intensities onto volume rendered surfaces clipped to that slice. The grey levels on the clipped slice can come from the intensities on which the surfaces are based or another registered data set. Figure 4 shows such a presentation via surface rendering on Pixel-Planes 4, but volume rendering seems to provide even greater capabilities. Volume rendering can more directly present these grey levels on the clipping slice either from the corresponding slice of original image intensities or capturing structures that range shallowly behind the clipping plane by computing the displayed levels from a projection, perpendicular to the slice, of intensities in a slab of specified depth beginning at that slice [Höhne, 1987]. In addition, with volume rendering the grey levels in the slab could be made to appear as grey levels in 3 space, using the volume rendering compositing to produce projections and cine display to provide rotation.

The volume rendering of grey scale values onto a clipping plane through an object surface involves generating two opacity and shade volumes, one for the object surfaces and one for the grey scale slice. In the object surface component opacity is set to zero on the viewer side of the clipping plane. In the slice component opacities must depend on a classification of which voxels are within the object image: full opacity must be used on the the region of the clipping plane within the object being displayed and zero opacity must be used outside the object, with a transition zone at the edges to avoid aliasing. For example, in figure 4 the opacity on the slice should be full inside the lung and zero outside the lung.

The above-mentioned techniques can be very effective in presenting relationships when one is combining a few opacities and shades into a single image. However, all too often the image appears to show a confusing fog when one attempts to combine numerous data sets. The choice of display and classification parameters, the design of opacity and shade functions of classification, and the design of means of compositing so as to provide effective visualization of numerous objects remains an area for research.

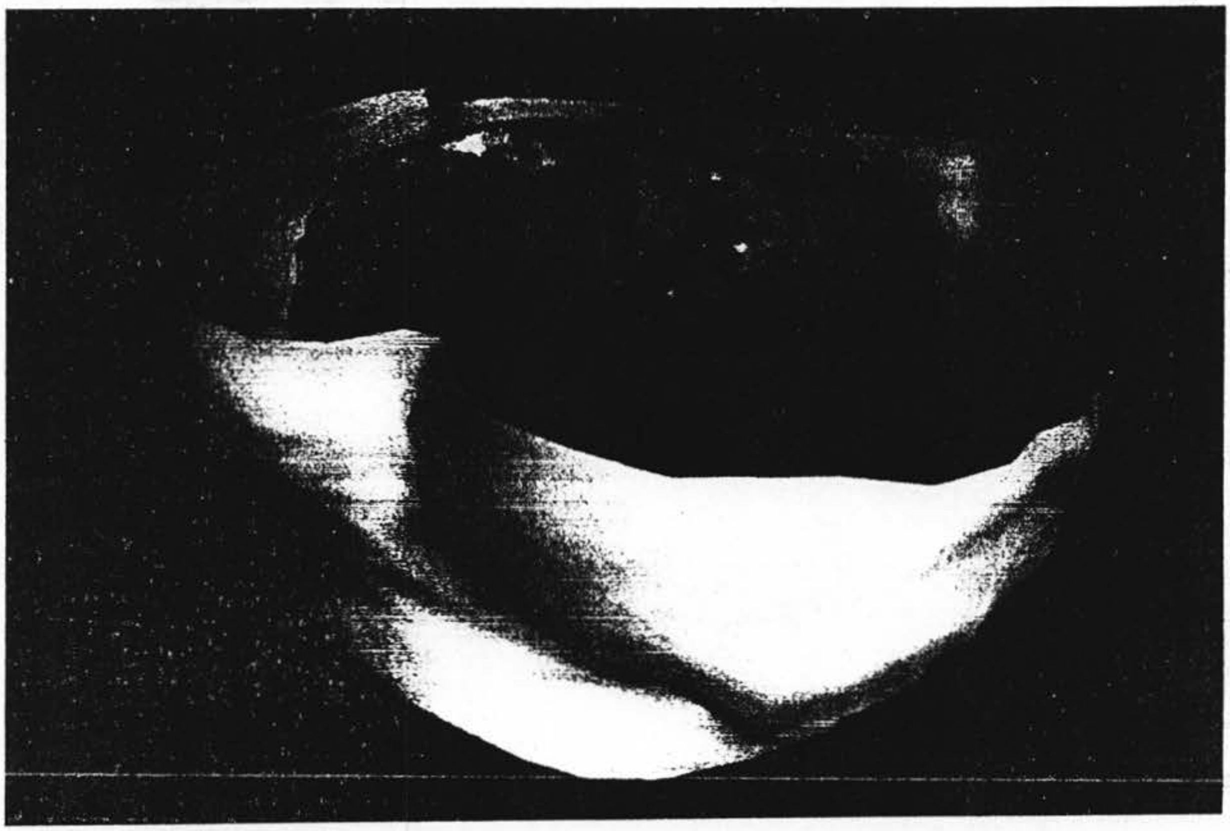

Figure 4. Lung perfusion from SPECT on a grey scale clipped slice through surface renderings of the body and lung surfaces from gamma CT.

4. SUMMARY

The most promising method of medical image display uses volume rendering to portray the relationship of many different kinds of objects: different anatomical objects defined directly from image intensities from one or more intensity data sets; treatment regions defined by the user; treatment objects defined by the user or derived from intensity distributions computed from these definitions; grey scale presentations. Our studies show great promise for such presentations. Besides the ever present challenge to produce better object classifications and exclusionary ROIs, the challenges are to provide means to define display parameters that will allow the simultaneous visualization of numerous objects and grey scale slices or slabs and to develop systems that will allow interactive specification of these display parameters with immediate image feedback.

5. ACKNOWLEDGEMENTS

Drs. Jordan Renner, Richard Davis, and Mark Schiebler have consulted on clinical aspects of imaging and display, and Mr. George Sherouse and Dr. Edward Chaney have consulted on and provided images from radiotherapy treatment planning. Images have been prepared and processing and display programs written by Mr. Andrew Skinner, Ms. Jiuqi Tang, Ms. J. Lynne Hendricks, Ms. Victoria Interrante, and Mr. Timothy Cullip. Dr. Benjamin Tsui's project on multimodality 3D display, the project of Dr. James Coggins, Mr. Timothy Cullip and Mr. Eric Fredericksen on interactive display of image region hierarchies, and the project on image description of Mr. John Gauch and Mr. Timothy Cullip have contributed significantly to the reported results. Finally, all of this work is made possible by our laboratories manager Dr. Graham Gash and by our administrative assistant, Ms. Carolyn Din, for whose help with this manuscript we are grateful.

The work leading to this paper was carried out with the support of NCI grant # P01 CA47982 and NSF grant # CDR 8622201.

6. REFERENCES

Drebin, R.A., Carpenter, L., and Hanrahan, P., "Volume Rendering," *Computer Graphics*, **22**(4):65-74, 1988.

Fuchs, H, Kedem, Z.M., Uselton, Naylor, B., "Optimal Surface Reconstruction from Planar Contours," *Comm. ACM*, **20**(10): 693-702, 1977.

Gauch, J., "Image Description via the Multiresolution Intensity Axis of Symmetry," *Proc. 2nd ICCV* , IEEE Catalog #88CH2664-1:269-274, 1988.

Höhne, K.H., DeLapaz, R.L., Bernstein, R., Taylor, R.C., "Combined Surface Display and Reformatting for the 3D-Analysis on Tomographic Data," *Invest. Radiol.*, **22**: 658-665, 1987.

Levin, D.N., Hu, X., Tan,, K.K., Galhotra, S., Chen, G.T.Y., Pelizzari, C. A., Balter, J., Beck, R. N., Chen, C., Cooper, M. D., "Integrated 3-D Display of MR, CT, and PET Images of the Brain," *Proc. NCGA*, vol 1:179-186, 1989.

Levoy, M., "Display of Surfaces from Volume Data," *IEEE Comp. Graph. & Appl.*, 8(3) :29-37, 1988a.

Levoy, M.,"Rendering Mixtures of Geometric and Volumetric Data," Technical Report 88-052, Comp. Sci., Univ. of North Carolina, 1988b.

Levoy, M., "Design for a Real-Time High-Quality Volume Rendering Workstation," Technical Report 89-010, Comp. Sci., Univ. of North Carolina (to appear in Proc. Chapel Hill Volume Visualization Workshop), 1989.

Mills, P.H., Fuchs, H., Pizer, S.M., Rosenman, J., "IMEX: a tool for image display and contour management in a windowing environment", *Proc. SPIE Medical Imaging III*, vol 1091, 1989.