Fuchs CAD 89 Berlin

3D Display with Minimal Predefinition

Stephen M. Pizer, Henry Fuchs, Marc Levoy,
Julian G. Rosenman, Richard E. Davis, Jordan B. Renner

Medical Image Display Research Group
University of North Carolina, Chapel Hill, NC, USA

Abstract
The portrayal of medical images in three dimensions can be strikingly improved in presentational usefulness and in user time requirements with methods based on volume rendering rather than surface rendering. We describe methods that lead to such improvement. These involve, in order of application, region-of-interest definition, as opposed to object definition; fuzzy surface specification; rendering not only fuzzy surfaces but together with them, surface models and grey-level slices; and fast display. These techniques lead to improved usefulness in many areas of medicine, including radiation oncology, surgery, multimodality image display, and diagnostic imaging from CT, MRI, SPECT, PET, and ultrasound.

Introduction

3D display of medical images has traditionally involved the predefinition of regions of the image volume corresponding to objects to be displayed, followed by the portrayal of the explicitly located surfaces of these regions. This surface rendering approach is flawed by requiring the user fully to comprehend and define the 3D objects to be portrayed before they are ever seen in 3D. The result is that object features can be missed and that only the highest contrast objects can be defined without hours of user time per 3D image.

Recently new methods for more directly displaying 3D data have been developed. Among the techniques developed have been the calculation of shading based on local intensity gradient [Höhne, 1986], the use of ray tracing to generate images [Schlusselberg, 1986], and improved algorithms for detecting surface voxels [Trousset, 1987]. More recently, our group [Levoy, 1988a] and others [Drebin, 1988; Sabella, 1988; Upson, 1988] have developed so-called volume rendering techniques, in which all voxels in the dataset are assigned a fuzzy surface likelihood and participate in the generation of each image. Such an approach has two major advantages over ones based on explicit

(binary) surface specification. First, minimal predefinition is required: only the carving out of a region excluding surrounding or adjacent objects that are likely to obscure or confuse the portrayal of the objects of interest. The surface likelihoods are directly calculated from the 3D intensity values in the original image data in the non-excluded region (call this the region-of-interest or ROI). Second, the presentation of the objects is improved since the method takes into account all the image intensities in a surface region rather than limiting the basis to binarily chosen surface points.

Surface rendering and grey-scale display techniques are not made obsolete by volume rendering. Surface rendering remains of significant use, in combination with volume rendering, for the portrayal of objects where binary specification is necessary: e.g., for the portrayal of prostheses, radiation treatment beams, isodose surfaces, surgical hazards, and other objects defined by the therapy planner, as well as anatomic objects whose volumes or surface properties must be measured. In addition, surface rendering may improve the appreciability of multiple, nested complicated objects that appear in human anatomy.

Grey-level display of the original slice intensities is not made obsolete because it portrays subtle intensity differences in a way impossible for any 3D volume display modality. Therefore, we have developed methods that allow the combined display of binary surfaces, surface likelihood, and superimposed grey-scale slices.

Volume rendering is most effective with fast display. Immediate display allows interactive selection of viewing parameters such as point of view, and allows ROIs and parameters affecting surface appearance to be chosen not when the image structures have been appreciated only in 2D slices but when this tuning seems necessary according to structure already portrayed in 3D. The effect of this tuning can be immediately appreciated in 3D, allowing real-time feedback to the parameter setting.

In the following we describe in order of their application each of the required steps, shown in Figure 1.

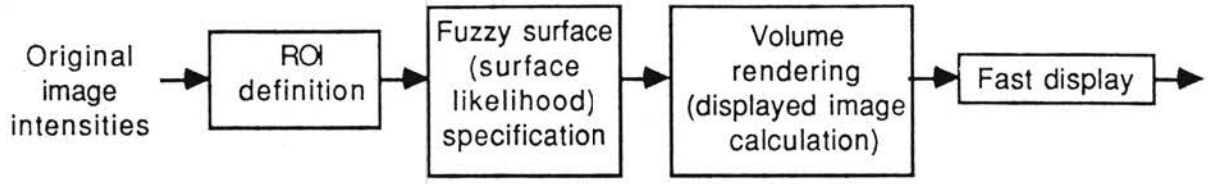

Figure 1. Display steps for volume rendering

Region-of-Interest Definition

Good display via volume rendering requires that surrounding and confusing regions be excluded before display. For example, to present the cortex of the brain from MRI (see Figure 2), the skull and its skin must be removed.

This selection of regions for exclusion, or for inclusion, is a far less demanding objective than the definition of the surface of the object to be displayed. Sometimes the ROI can be specified by a simple geometric construct such as a rectangular prism or an ellipsoid. In other cases simple sculpting may be necessary. Both of these approaches require the visualization of the selected region relative to anatomy during the ROI selection, either in 3D or on an array of grey-scale slices of the original image data. Interactive display is thus necessary.

An alternative to specification of the ROI via simple geometric shapes is based on automatically precomputed image descriptions that divide the image into "sensible", coherent regions, probably hierarchically by scale [Pizer, 1989]. The user can then point to a voxel (perhaps displayed as a pixel on a grey-scale slice) and cause a whole image region to be selected and displayed [Coggins, 1989]. This region may frequently be used without editing in selecting or excluding voxels for volume rendering. Gauch [1988] has recently demonstrated the feasibility of computing such regions in 2D images based on a construct called the *Intensity Axis of Symmetry* that defines regions in terms of symmetry relations on the image intensities in an

intensity "ridge" or "course" (see Figure 3). Cullip [1989] is now extending this approach to 3D.

Fuzzy Surface Specification

Volume rendering computes opacities and shades from a surface likelihood value at each voxel. Such a value is far easier to specify with adequate accuracy than binary surface locations, because the display method is more forgiving to errors in this non-binary value.

Drebin [1988] and Levoy [1988a] have described methods for using intensity value and intensity gradient to calculate a surface likelihood or "strength" for every voxel in the original image data. The basic idea is to let the surface strength increase both with the closeness of the voxel's intensity to some intensity range associated with the object to be displayed, and with the magnitude of the intensity gradient at the voxel. It seems possible to have the surface strength depend also on higher order local geometric properties such as curvature, global geometric properties such as symmetry or surface continuity, or surface probabilities calculated using knowledge of the imaging process [Lin, 1989].

Binary surface specification is handled by volume rendering as a special case, where the probabilities are zero or one. This is useful in the portrayal of treatment objects (see Figure 4) and of scalar intensity distributions without edges, such as radiation dose distributions that need to be visualized relative to anatomy in radiotherapy treatment planning. Surface likelihoods of unity can be set at locations having selected levels of the dose. Binary surface specification will be also be useful for portraying anatomic objects whose edge contrast is too low or diffuse to allow surfaces to be portrayed by direct volume rendering. In this case, the user can define object surfaces by the standard manual techniques, and the resulting surface tiles can be inserted into the volume rendering.

Volume rendering, as described by Levoy and Drebin, also requires a surface normal for each voxel. This is normally taken to be the

direction of the intensity gradient, though higher order measures of local shape could also have an effect.

Volume Rendering

In volume rendering opacities and shades are computed from the surface likelihood at each voxel, and rays are cast to calculate the effect of the shades as attenuated by the opacities on the display pixel through which the ray passes. Our results show the importance of careful anti-aliasing in the interpolations involved in computing shades and opacities along the rays from their values at the discrete voxel centers (see Figures 5 and 6). They also show the usefulness of specular highlights, fixed light positions and rotating objects, and color use (see Figure 4) to distinguish object types.

Levoy has developed techniques [Levoy, 1988b] for converting surfaces defined by polygonal tiles into the opacities and shades used by volume rendering and doing so in a way that avoids aliasing artifacts. Multiple surfaces, some rendered directly by volume rendering and others by polygonally modeled surfaces rendered into the surface likelihood values, can be used with color and transparency to allow the understanding of multiple, complicated objects, as, for example, are found in the inner ear (see Figure 6). It is frequently useful to portray grey values superimposed on 3D images, with the 3D objects providing context by which to understand the grey values. Höhne [1987] has shown the effectiveness of this approach with grey-scale values coming from the original intensity data on which the 3D display is based. In our group Tsui [1988] has shown that this idea is also effective for showing physiological function from one 3D image intensity set in the context of anatomy coming from another. Our approach is to render such intensities onto planes which are combined with volume rendering via texture mapping onto planes that clip the volume rendered 3D region [Levoy, 1988b].

Fast Display of Volume-rendered Images

Volume rendering of 256 X 256 X 256 image intensities can be accomplished in approximately 10 minutes of computer time on a Sun 4 workstation, using techniques developed by Levoy [1988c, 1989]. These

techniques allow the re-rendering from a new point of view in only 2 minutes. Furthermore, progressive refinements are possible which allow a coarse rendering in only a fraction of this time, with improvement to full quality occurring over a few minutes when interaction stops.

The techniques described above lead to an even greater promise for volume rendering: allowing real time exploration, in 3D, of 3D intensity data. On machines such as Pixel-Planes 5, a parallel graphics and image processing engine being developed in our laboratory [Fuchs, 1989] we calculate that it will be possible to produce coarse renderings from 256^3 voxels at the rate of 20 per second. Thus it should be possible to vary parameters of the voxel classifications (such as ROI or the opacity function's slope vs. gradient) or viewing parameters (such as point of view, color, transparency, or texture-mapping clipping plane) and see the result fast enough for dynamic adjustment. This approach promises to allow the same sort of tuning of the image to optimally portray an anatomic region that is presently achieved by interactive intensity windowing with 2D images. The power of this ability to explore directly in 3D will not be fully known until this ability is realized, but we think it may qualitatively increase the usefulness of 3D medical image display for diagnostic purposes and improve the already accepted uses of 3D display in planning for therapy in radiation oncology and surgery.

Clinical Results

This section characterizes the clinical effectiveness of 3D display based on volume rendering, as illustrated by the images already shown and one new image. Illustrated properties of volume rendering are also discussed.

Figure 2, showing a volume rendering of the cortex from MRI, illustrates that a complicated surface can be well portrayed without much manipulation. Only the image regions corresponding to the skin on the skull were excluded. It was not necessary to remove bony skull since its intensity is very different from the cortex. Also, since a single layer of surface elements can be made to produce only partial transparency, it was not necessary to remove the closely fitting but

thin arachnoid mater to visualize the cortex, which is thicker and thus builds up the opacity that makes it appear as a surface. With such an image of a complex surface, we have found improved visualization by putting the brain into rotation via a cine loop, especially if the light source has a fixed position.

Figure 4 illustrates the potential of volume rendering in radiation treatment planning. It shows the relation between a radiation beam, in blue, and a patient's brain tumor and head. It is our experience that without high quality 3D graphics it is difficult for a clinician to understand the complex spatial relationship between the radiation beam positions with their associated dose distribution, the target treatment region, and radiosensitive normal tissue in a fully 3D radiation treatment plan. The figure illustrates how combination of objects represented by surface likelihood and by polygonal tiles may increase the clinician's comprehension of a given treatment plan. Because of its inherent detail, volume rendering provides an excellent understanding of the patient's anatomy, represented by surface likelihoods computed from CT image data; also, the method's ability to use color allows the clinician to distinguish different tissue types, such as the flesh in red and the bone in yellow. At the same time the method allows the appreciation of the user-defined (and thus tiled) radiation beam and target region, the latter consisting of both the imaged tumor and surrounding tissues into which the tumor may have spread. The result makes it easy to see, for example, if the entire tumor target is covered by the beam or if the beam is unnecessarily impacting upon normal tissues such as the patient's eye.

The clinician typically must interactively select from or vary radiation beam arrangements. Choosing arrangements that correctly intersect or avoid anatomic regions requires volume renderings that are produced fast enough for the viewpoint to be dynamically varied. The final evaluation, requiring the clinician to understand the relation of the radiation dose distribution and the patient's anatomy, can be made from the combination of transparent, colored isodose surfaces, volume rendered directly from the dose distribution, and the anatomy and tumor target, volume rendered directly from the CT image

data. A more detailed appreciation of the fine details of the radiation dose distribution might well require the display of (not necessarily transverse) slices of the CT data, in grey scale with colored isodose contours, superimposed for context on the volume-rendered anatomy clipped by the slice.

Figure 5 and 6 illustrate the use of volume rendering for surgical planning. Figure 5 shows a volume rendering, based on CT data, of a severe fracture of the os calcis, the largest bone in the heel. The ability and decision to repair such a fracture operatively depends both upon the extent of involvement of the joint surfaces of the os calcis and upon the size and positional interrelationships of the various bone fragments. This information can be gleaned from conventional tomography or from computed tomography, with normal slice display or surface rendering, but the exquisite detail afforded by volume rendering permits a more precise understanding of the anatomy of the fracture. Moreover, surface renderings of bone structures may incorrectly portray holes where there is thin or osteoporotic bone. The figure illustrates how volume rendering avoids this problem -- even flakes of bone fragments appear as solid structures, and thin bone appears partially opaque rather than as a hole.

The use of 3D reconstructions in orthopedic, craniofacial, and oral surgery has been well established [Moaddab, 1985; Vannier, 1985; Fishman, 1988], but widespread clinical acceptance has been slow, perhaps because of its inability to fully capture subtle skeletal detail. The superior image quality of volume rendering, limited only by the resolution of the original CT data, provides useful information for those surgical disciplines which may require greater precision in the evaluation of skeletal anatomy, such as otorhinolaryngology -- head and neck surgery. Visualizations of the human temporal bone (see Figure 6), which encases numerous tiny vital structures in complex interrelation and close approximation, may prevent surgical errors in special clinical circumstances such as congenital malformations, prior temporal bone surgery, or altered temporal bone anatomy due to trauma or neoplasia. Preliminary studies suggest that both radiation exposure

and risks of sedation will be equivalent to those of conventional 2D CT scanning and thus are justifiable in terms of patient safety.

The pictures in figure 6, based on 20 axial CT slices with 1.5 mm spacing, show the left temporal bone of a 51 year old woman scheduled to undergo revision ear surgery for recurrent temporal bone disease (chronic mastoiditis). Cine loops, created by compiling 60 similar images by rotating through a 180 degree arc at 3 degree increments provide improved 3D comprehension and allow inspection from any component perspective (Figure 6a). A different 60 frame movie, in which the temporal bone (shown from the surgical perspective) is progressively "sanded away", one pixel-thick layer at a time, to reveal hidden underlying anatomy, has also proved useful. The viewer can stop at any point for closer scrutiny of a particular frame (Figure 6b).

Figure 7 demonstrates volume rendered images of a CT scan of a normal volunteer's kidneys. It shows that volume rendering can provide excellent visualization of low-contrast soft tissue structures, here especially the major abdominal blood vessels and those vessels supplying the kidneys. However, the kidneys appear distorted due to a respiratory motion artifact -- volume rendering's sensitivity to motion artifacts make research into special treatment of these artifacts necessary.

The early results of ourselves and others [McCann 1988] suggest that volume rendering has considerable potential for the display of the 3D pulse-echo ultrasound images that are beginning to be experimented with [von Ramm, 1988; McCann, 1988; Hottier, 1988]. Here surface likelihood is not to be measured by intensity gradient, since the signal itself measures a distorted form of surface strength. However, it is possible here to take account of image continuities in the time dimension in computing surface likelihood, and in the long term it may be possible to display real-time volume renderings as the 3D ultrasound images are collected.

Summary and Conclusions

Volume rendering can provide a number of benefits that will strongly aid clinical treatment planning, diagnosis, and followup. Volume renderings of objects that appear in image data with high contrast show more detail and more accurate detail than surface renderings. Objects that appear less sharply in image data can provide adequate 3D comprehension while avoiding the enormous time consumption of object definition that is the rate-limiting process with clinical application of surface rendering techniques. In imaging areas like nuclear medicine and ultrasound, in which surfaces are inherently fuzzy, volume rendering may be useful where surface rendering was close to useless. Especially important for diagnosis, volume rendering may allow exploration in 3D of the image data to find out what 3D structures and relationships the data show. Clinically important uses of this technique in surgery planning, radiotherapy treatment planning, and medical imaging diagnosis abound.

However, to obtain these benefits, much work is still needed. Better classification techniques for defining ROIs and the surface likelihood of selected organs are necessary. Improvements in mapping surface likelihood to opacity will improve image quality when surfaces are especially fuzzy. Methods for avoiding the effects of motion or correcting for motion are very important to avoid the major artifacts that motion now causes for volume rendering. Methods for real-time rendering of precollected images must be developed not only to add convenience to preparation of images for therapy planning but to allow at all the explorational forms of 3D display that most diagnostic uses await. And methods of real-time rendering of real-time image data will allow volume rendering to be used to visualize dynamic processes in the body while they happen.

Acknowledgements

Dr. Mark Schiebler has consulted on clinical aspects of MRI 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, 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. We are also indebted to NeuroCT technicians Bonnie Gordon and Pamela Buckingham, who have spent considerable time archiving CT data sets to be used for 3D display. 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.

We much appreciate 3D MRI images provided by Mr. Hartmut König of Siemens, Erlangen. The work leading to this paper was carried out with the support of NCI grant # P01 CA47982 and NSF grant # CDR 8622201.

References

Coggins, J., Cullip, T., Fredericksen, E., Pizer, S., "A Data Structure for Image Region Hierarchies," Internal Report, Department of Computer Science, University of North Carolina, 1989.

Cullip, T., "Intensity Axis of Symmetry and its Application to Image Segmentation and Hierachy Manipulation", *Radiology Research Review*, Univ. of NC Dept. of Radiology, 1989.

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

Fuchs, H., Poulton, J., Eyles, J., Greer, T., Goldfeather, J., Ellsworth,D., Molnar, S., Turk., Tebbs, & Israel, L., "A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," Technical Report 89-005, Computer Science Department, University of North Carolina at Chapel Hill (submitted for publication), 1989.

Fishman, E.K., Magid, D. Ney, D.R., Kuhlmand, J.E., Brooker, A.F., "Three Dimensional Imaging In Orthopedics: State of the Art 1988", *Orthopedics*, **11**(7):1021-1026, 1988.

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

Höhne, K.H. and Bernstein, R.,"Shading 3D-Images from CT Using Gray-Level Gradients," *IEEE Transactions on Medical Imaging*, **5**(1):45-47, 1986.

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

Hottier, F., Personal Communication, 1988.

Levoy, M, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*, **8**(3):29-37, 1988a.

Levoy, M., "Rendering Mixtures of Geometric and Volumetric Data," Technical Report 88-052, Computer Science Department, University of North Carolina at Chapel Hill (submitted for publication), 1988b.

Levoy, M., "Efficient Ray Tracing of Volume Data", Technical Report 88-029, Computer Science Department, University of North Carolina at Chapel Hill(in review, *ACM Transactions on Graphics)*, 1988c.
Levoy, M., "Volume Rendering by Adaptive Refinement," *The Visual Computer*, **5**(3), 1989a.
Lin, W., *Probablistic Surface Definition in 3D Ultrasound Images*, Ph.D. Dissertation Proposal, UNC Department of Computer Science, 1989.
McCann, H.A., Sharp, J.C., Kinter, P.M., McEwan, C.N., Barillot, C., Greenleaf, J., "Multidimension in Ultrasound Imaging for Cardiology," *Proc. IEEE*, **76**(9):1063-1073, 1988.
Moaddab, M.B., Dumas, A.L., Chavoor, A.G., Neff, P.A., Homayoun, N., "Temperomandibular Joint: Computed Tomographic Three-Dimensional Reconstruction," *Am. J. Orthod. Dentofacial Orthop.*, **88**(4):342-352, 1985.
Pizer, S.M., "Multiscale Methods and the Segmentation of Medical Images," to appear in *The Formation, Handling and Evaluation of Medical Images*, Springer-Verlag, Berlin, 1989.
Sabella, P., "A Rendering Algorithm for Visualizing 3D Scalar Fields," *Computer Graphics*, **22**(4):1-58, 1988.
Schlusselberg, D.S. and Smith, W.K., "Three-Dimensional Display of Medical Image Volumes," *Proc. NCGA* '86, Anaheim, CA, **3**:114-123, 1986.
Trousset, Y. and Schmitt, F.," Active-Ray Tracing for 3D Medical Imaging," *Proc. EUROGRAPHICS* '87: 139-149, 1987.
Tsui, B.M.W., Perry, J.R., Gilland, D.R., Interrante, V.L., Fuchs, H., & McCartney, W.H., "Preliminary Study of Simultaneous Three-Dimensional Display of Anatomical Functional Information Obtained with SPECT," *Proc. Soc. Nuclear Medicine*, 1988.
Upson, C. and Keeler, M.,"VBUFFER: Visible Volume Rendering," *Computer Graphics*, **22**(4):59-64, 1988.
Vannier, M.W., Conroy, G.C., Marsh, J.L., Knapp, R.H., "Three - Dimensional Cranial Surface Reconstruction Using High-Resolution Computed Tomography," *Am. J. Phys. Anthropol.*, **67**:299-311, 1985.
von Ramm, O., Smith, S. W., Sheikh, K.H., Kisslo, J., "Real-Time 3-Dimensional Echo Cardiography", presented in *Circulation*, Supplement II, **78**:35, 1988.

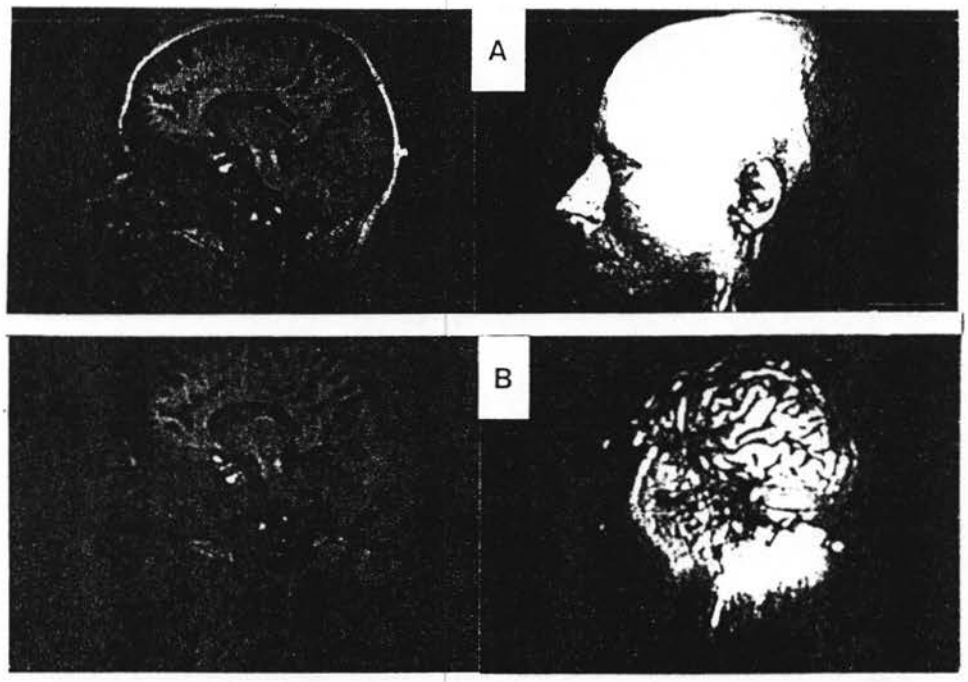

Fig.2a. MRI Sample slice and volume rendering of skin surface (MRI data courtesy of Siemans).

Fig.2b. MRI sample slice edited to exclude skin region, and subsequent volume rendering of exposed cerebral cortex.

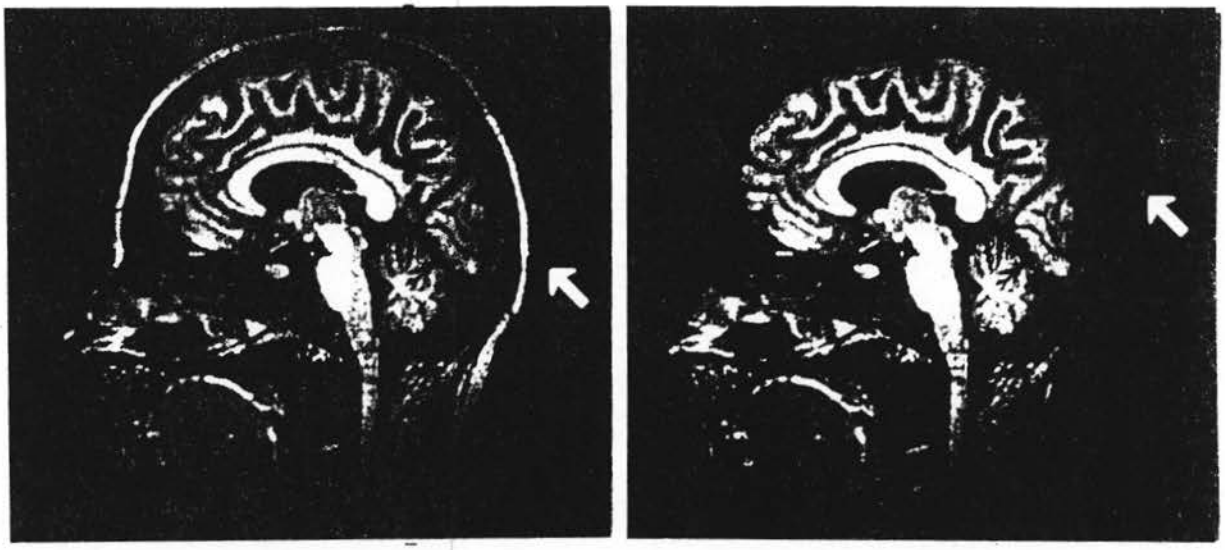

Fig.3. Sensible region computed via Intensity Axis of Symmetry of 3D MRI data (MRI data courtesy of Siemens).

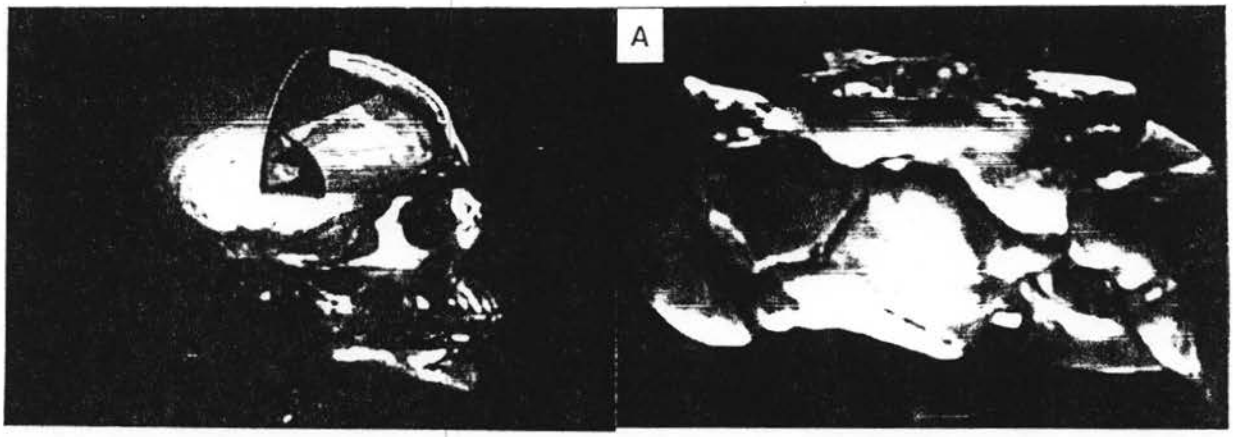

Fig.4. Polygonally tiled radiation treatment beam and tumor target in volume rendered head from CT.

Fig.5a.

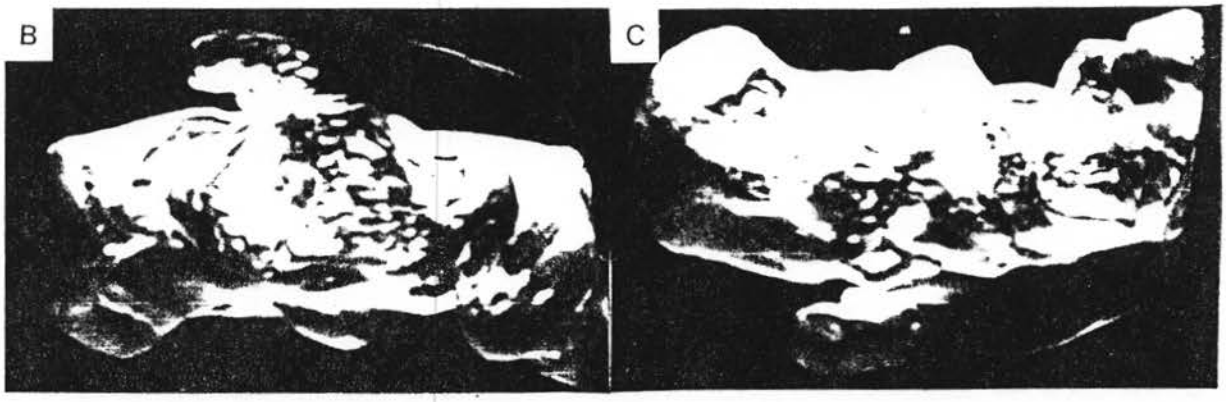

Fig.5. Volume rendered medial (5a), superior (5b), and inferior views (5c) of fractured human heel from CT.

Fig.6. Volume renderings for surgical planning of temporal bone from CT data. Fig.6a shows representative frames from a rotational cine loop. A 3.0 cm slice of the skull is seen from a left superior oblique viewpoint. The left zygomatic arch(z), posterior nasal septum(s) and posterior cranial fossa(f) are indicated. Fig.6b1 shows the left temporal bone from the surgeon's perspective. The patient is lying face up and is viewed down the external ear canal(e). In fig. 6b2, the temporal bone has been "sanded away" to reveal the cochlea(c), the internal auditory canal(i), the contralateral internal auditory canal(a) and the contralateral sigmoid sinus (ss).

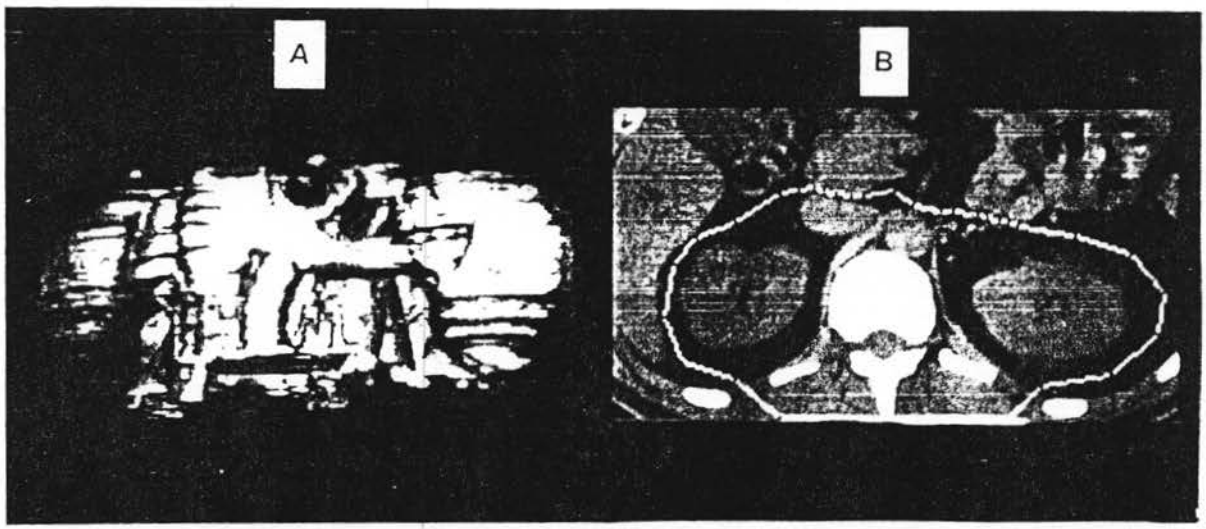

Fig.7. Volume rendering of kidneys from CT DATA (7a), with one slice (7b) showing included ROI.