Three-Dimensional Image Display in Medicine

Nicholas J. Mankovich, Douglas R. Robertson, and Andrew M. Cheeseman

This article is a tutorial on the methods used to create three-dimensional (3-D) images for use in displaying patient anatomy. This new view into anatomy has developed over the last 10 years from the need of surgeons, radiation therapists, and radiologists to integrate the many images resulting from the recent growth in tomographic imaging including computed tomography (CT) and magnetic resonance imaging (MRI). CT and MRI studies result in 30 to 100 images. 3-D imaging processes and integrates this image data volume and extracts more meaningful, derivative images via multiplanar reconstruction (MPR), shaded surface processing, or volumetric processing. MPR reslices the image volume to produce novel views of patient anatomy while retaining the image voxel intensities. Realistic shaded surface display of 3-D objects can involve extensive processing of the images to create computer representations of objects rendered into a displayable 3-D scene. Volumetric imaging combines the voxel processing of MPR with the techniques of tissue classification and surface shading to produce novel projections of the image data volume that allow automated creation of 3-D scenes without recourse to the complexities of object delineation. As the ultimate 3-D display, recent advances in computer-aided design (CAD) and computer-aided manufacturing (CAM) allow the fabrication of physical models of anatomy using computer-controlled milling machines. New technology that actually builds the model layer by layer from a liquid plastic offers the possibility of complete models with intact internal anatomy. The growth in 3-D is certain as hardware and software costs decrease and medical professionals find further applications for this technology.

© 1990 by W.B. Saunders Company.

KEY WORDS: three-dimensional imaging, image display, shaded surface display, multiplanar reconstruction, computer graphics, computed tomography, magnetic resonance imaging.

THE GROWTH of medical computer technology has brought change to all aspects of medicine, particularly in diagnostic radiology, where imaging techniques allow internal anatomy to be visualized in a slice-by-slice (tomographic) fashion. Attending physicians, sur-

geons, radiation therapists, and patients all rely on the radiologist's expert interpretation of computed tomography (CT), magnetic resonance imaging (MRI), nuclear medicine, and ultrasound (US) images for their assessment of patient anatomy and pathology.

The surgeon and radiation therapist have needs that transcend diagnostic interpretation: they must carefully plan their approaches and techniques preoperatively.¹ This planning requires a detailed knowledge of particular patient anatomy and extensive clinical experience in specialized surgery. Fortunately, the growth of tomographic imaging in medicine has been paralleled by the emergence of powerful three-dimensional (3-D) computer graphic tools that are becoming a part of medical imaging.²⁻⁴ It is possible, and now clinically practical, to use computer hardware and software to convert the 20 to 60 sectional images from CT or MRI into realistic pictures of patients' bones, muscles, and organs.⁵⁻¹³ These 3-D images and, more recently, computer-generated anatomical models are used by the surgeon for preoperative planning, postoperative assessments, and patient education.

By far the most common source of 3-D information in medicine is the tomographic scanning of some portion of the human body (Table 1). CT rotates an x-ray fan beam generator around the patient and then picks up differential x-ray attenuation with a bank of detectors on the opposite side of the patient. A large number of these beam attenuation profiles are then recon-

From the Department of Radiological Sciences, University of California, Los Angeles, and the Department of Radiology, Johns Hopkins University, Baltimore, MD.

Address reprint requests to Nicholas J. Mankovich, PhD, Department of Radiology, Olive View Medical Center, UCLA, 14445 Olive View Dr, Sylmar, CA 91342.

© 1990 by W.B. Saunders Company.
0897-1889/90/0302-0005\$03.00/0

Table 1. Sources for Three-Dimensional Images

Tomographic images (slices)	
Anatomical	
	Computed tomography (CT)
	Magnetic resonance imaging (MRI)
	Ultrasound scanning (US)
	Other: conventional tomography, histology, etc
Functional	
	Positron emission tomography (PET)
	Other nuclear medicine procedures (SPECT, etc)
Microscopic	
	Limited focal-plane (using focus as slice localizer)
	X-ray microtomography
Projection—limited planar views	
	Biplane angiography
	Dynamic spatial reconstructor (Mayo Clinic)
	Stereo pairs: angiography/mammography
Surface	
	Pattern illumination via shading, moire, or laser
	Sonic detectors
Other	
	Molecular modeling
	General data visualization

structured into an image representing a cross-sectional tissue density map showing the patient anatomy in the region scanned. A typical CT image has 4,096 possible gray levels (12 bits) in a 512×512 matrix of picture elements (pixels). Within the sectional image, the square pixels range from 0.5 to 1.0 mm on a side with slice thickness from a 1.0 to 10.0 mm, according to the width of the fan beam. The addition of this third dimension, slice thickness, changes the picture element into a volume element (voxel). Table 2 lists the most salient features of CT and MRI along with some typical operational parameters.

A typical 3-D CT study performed with 2-mm contiguous slices results in a patient skin dose of

6 rad in the middle of the area scanned. This can be lowered to 2 to 3 rad¹⁴ if noncontiguous scans are used and the tube current (in milliampereseconds) is reduced. The lowered-dose images suffer some loss of soft tissue detail but generally maintain an excellent separation of bone and soft tissue. However, if the goal of 3-D CT is the delineation of soft tissue structures, then dose reduction is difficult.¹³

An attractive alternative to the x-ray radiation of a CT scan is the MRI scan, which uses a combination of magnetic and radio frequency gradients to collect information on the distribution of protons in the living subject. This technique uses no ionizing radiation; thus, the biological effects on the subject are minimal. However, the protocol for scanning and subsequent image processing can be complex.

3-D MRI is limited by nonuniform signal intensity and partial volume averaging due to its relatively low signal to noise characteristics. The gain is in MRI's increased sensitivity to subtle chemical differences among tissues. The net effect is that MRI can produce tomographic images that show a much clearer delineation among both diseased and healthy soft tissues. Thus, MRI has great potential in 3-D imaging, although the limitations of its low signal-to-noise ratio and larger slice thickness present challenges to the development of 3-D MRI.^{12,15-17}

There are many other ways of acquiring data for 3-D processing although none of them have reached the popularity of CT. US has been used in cardiac studies to estimate ventricular volume.¹⁸ Biplane angiography helps in localizing pathology in both cardiac and cerebral studies.^{19,20}

Table 2. Characteristics of Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) as They Apply to the Three-Dimensional Reconstruction of Image Data

Characteristic	CT	MRI
Matrix size	512×512	256×256
Density resolution	12 bits (4,096 levels)	16 bits (actual 7 bits = 128 levels)
Pixel dimension	0.5×0.5 mm	0.6×0.6 mm
Slice thickness	1.5 to 2.0 mm	2.0 to 3.0 mm
Slice spacing	2 mm ranging from 1 to 2 mm	2+ mm
Voxel dimensions	$0.5 \times 0.5 \times 2.0$ mm	$0.6 \times 0.6 \times 2.0$ mm
No. of slices/3-D study	30 to 70	20 to 128
Study protocol	Simple: choose location, technique, thickness, spacing	Complex: choose location, thickness spacing, parameters (TE, TR, . . .), and interleave
Patient risk	High radiation dose (4 to 6+ rad)	Low, no ionizing radiation
Data storage	8 to 18 million voxels	1 to 8 million voxels

NOTE. Values are presented for the typical CT or MR scanner, exact specifications vary.

The familiar fan-shaped image of US demonstrates its fundamental tomographic nature. Once the probe location is exactly known, a series of images can be collected that represent slices of the object of interest. Numerous angled slices may be taken from the same probe location or from precisely spaced parallel planes. Biplane angiography relies on the known projection geometry to back-project reference points into 3-D space. This triangulation method can be used to trace the course of vessels in the coronary artery tree or the cerebral arteries. Of course, the (x,y,z) location of an object derived from only two views is much less precise than the hundreds of views afforded by the multiprofile technique of CT.

Other methods restricted to surfaces exist for the acquisition of 3-D information. These are of less interest to the neurosurgeon or orthopedic surgeon than to the plastic or maxillofacial surgeon. In these professions, the external appearance of the patient can be primary in planning the course of surgery and predicting its outcome. Various methods have been devised that automatically give surface contours of a pattern-illuminated object. Originally, twin projectors were used with grid patterns illuminating a central object.^{21,22} Video images of the illuminated object were then computer processed to yield information about the distance of object surface points from the camera. More recently a scanning laser has been used to accurately illuminate an object while a digitizing camera scans the object. This technology has been successfully used in the creation of computer-machined sculptures,²³ but overall medical application is limited by its restriction to surface features.

PROCESSING FOR 3-D DISPLAY

Successful 3-D displays provide the viewer with the perception of the object's shape, orientation, relative location, relative size, and possibly the internal structure. Although it is beyond the scope of this discussion, this display can integrate anatomical and functional data (eg, from positron-emission tomography [PET]), allowing the accurate planning necessary for stereotactic surgery.^{10,24,25} When considering the basic problem of viewing on a two-dimensional (2-D) display device, a 3-D scene might be rendered as surface contour lines (vector graphics) or realistic-

looking shaded surfaces (shaded raster graphics) composed of one or more opaque or transparent objects with realistic shadows from multiple light sources. To enhance the 3-D impression the scene might exaggerate the perspective, displaying distant parts of objects smaller than equally sized nearer objects. The scene might contain all of the subtle cues that we see in real scenes such as reflection, highlights, texture, and refraction through a transparent object. Further, all of this might be combined into a short moving-object sequence (movie loop) showing a rotating set of objects in which we can see parts appear and disappear as our viewpoint changes. At their extreme, if image acquisition devices were of ultrahigh-resolution, then the 3-D displays could produce images indiscernible from a photograph of the physical object (photorealism). Ultimately, 3-D displays are limited by the spatial and density resolution of patient scanning hardware.

The key elements in processing for 3-D image display include finding the surface, determining the normal vector at each location, and applying an efficient and realistic rendering.²⁶ There are many ways to classify the processing of data for 3-D reconstruction. In this discussion, a distinction is drawn between voxel-based and object-based processing, each relying on a different method of translation from the many sectional slices comprising the image volume into a new view into the 3-D structure. To a large extent, the method chosen determines the amount of computer and operator processing required for 3-D display.

The starting point for 3-D display of medical tomographic information is the volume element, or voxel. Each of these elements in a given slice represents an average of tissue characteristics found in that small volume. Voxel-based processing makes determinations about each voxel and decides to what degree each should contribute to the final 3-D display. Object-based processing uses voxel information to transform the images into a collection of objects with subsequent processing concentrating on the display of the objects.

Each image acquisition modality possesses limitations in detection and image reconstruction. Variability is present in all systems and each voxel is a statistical sample that, in the

context of subsequent 3-D imaging, can benefit from processing to mitigate the effects of noise and finite sampling limitations (partial volume effects). Such neighborhood processing includes median filtering, regional averaging, or any other image processing technique that attempts to enhance structures of interest. An inherent problem in most tomographic scanners is the non-cubic voxel. For subsequent 3-D processing, it is frequently useful to change from the rectangular voxel (eg, CT $.5 \times .5 \times 1.5$ mm) to a cubic voxel by applying image interpolation techniques.^{12,27-29,30} Thus an image acquisition of 40 slices at 1.5 mm thickness might be interpolated into an isotropic patient data volume of 120 slices, each with cubic $0.5 \times 0.5 \times 0.5$ mm voxels. Our discussion will move from the simplest voxel-based method of 3-D reconstruction—multiplanar reconstruction—to object-based reconstruction, and finally return to the more complex voxel-based method—volumetric imaging.

Multiplanar Reconstruction

The simplest way of processing the data volume is multiplanar reconstruction (MPR) in which the image volume is interpolated and resliced to produce novel views of patient anatomy.^{14,31,32} Although not strictly a 3-D display, MPR has come to be treated as an added processing step available to the physician to improve understanding of true 3-D anatomy. Dating back to 1975,³³ this voxel-based processing retains the original voxel intensity information. In the past few years MPR has become available on both CT consoles and in stand-alone 3-D workstations. In its simplest form, MPR can be used to automatically generate coronal or sagittal sections from a single transaxial image series. When oblique sections are required, the multislice tomographic data are first reprojected into an anterior-posterior or lateral image resembling a conventional x-ray. The operator uses a trackball or joystick to position a line on the display screen until it overlies the desired region of anatomy. The computer then resamples the interpolated data volume to provide a slice image along that plane. More complex curving lines are used to produce images that represent curved surfaces through the data volume. One example of this technique is in a CT spine series, where the reconstruction program is used to produce a

series of slices parallel to the spinal canal, thus demonstrating the course of the spinal nerves or bony abnormalities.

The initial reprojection of the image used in locating the new plane of interest in MPR has proven to be a useful concept. Because each voxel in a CT study represents the x-ray attenuation of that particular pixel, we can sum all pixels along any proposed x-ray path.³⁴ The total attenuation at the ray exit point is identical to the x-ray intensity that would be found if the image were formed as a conventional projection radiograph. This method, when combined with tissue classification and various voxel neighborhood operations, has tremendous potential in viewing 3-D structures. This will be discussed later in the volumetric imaging section.

Voxel-based processing for reprojection and multiplanar imaging is relatively simple. In the case where orthogonal projections are sufficient, the problem reduces to one of data volume addressing. Usually, this three-view display is not sufficient, and volume addressing must be accompanied by linear or cubic interpolation because the slicing plane intersects multiple voxels. This interpolation yields 3-D images that have spatial detail superior to the simple x , y , or z axis projections at the cost of more extensive computation time. Currently, most manufacturers of CT scanners offer a software package designed to perform multiplanar reconstruction.

Object-Based Processing

To the orthopedic surgeon, it is often sufficient to understand only the relative positioning of the bones. To accomplish this, object-based display methods convert a complex 3-D image data set of many millions of voxels to a few pictures of bony anatomy. Image processing determines which voxels are part of the object and should be displayed, and which are not and should be discarded. This process of identifying components of an image is known as image segmentation. In the simplest bony object delineation, the application of an intensity threshold eliminates all nonbone objects. Once the image set has been put through this binary filter (there/not there) a variety of algorithms are used to produce a picture of the anatomy with realistic object depiction. This transformation of data from an internal object representation to a viewable dis-

Fig 1. The general flow of information and processing in the creation of a 3-D object-based display.

play (sometimes called a scene) is called *rendering*. Figure 1 diagrammatically shows the process of 3-D image construction from original object including the steps of scanning, data filtering, object delineation, and rendering.

Data filtering. In its simplest form, image segmentation invokes a binary decision process that interrogates each voxel to determine its inclusion in the object set. In medical images this is most frequently done on CT images to delineate the bones. The CT operator indicates the CT number threshold, say 200, above which the voxel will be considered bone, all subthreshold voxels are considered nonbone and hence become empty background. As structures become more difficult to delineate (as in soft tissue in CT or MRI), more sophisticated image processing is used to enhance the data before, during, and after the binary decision process. Because of the critical nature of this first step in determining the inclusion or exclusion of objects, a variety of techniques are available beyond simple thresholding.^{12,17,35-38}

Object delineation. In the simplest case, once a threshold has been specified, the image data volume, now containing only object voxels, is processed to locate all objects and change their representation from image voxels into an object model. Because there are many ways of presenting objects in displays and representing them in computers (surfaces, vectors, points, etc), there are many algorithms for going from threshold to an object model suitable for rendering. Basically, the image to object transformation can be accomplished by either 2-D or 3-D processing. 2-D processing takes each thresholded image slice, in turn, and identifies appropriate object voxels, ignoring adjacent slices. The exact type of 2-D processing depends on the desired result. Boundary extraction methods rely on finding all voxels that are adjacent to nonobject (background) voxels. The resulting data set is the location of all voxels that comprise the outer (or inner) border (surface) of the object. This is often referred to as the object contour. Volume extraction methods

rely on locating all voxels that actually comprise the object (including its surface).

There are many ways to perform boundary extraction. A simple 2-D slice method called boundary following starts with an operator (or algorithm) who locates a starting point at the surface of an object. Figure 2 shows how the border is found by examining each pixel in a 3×3 (9 pixel) neighborhood around the starting pixel.³⁹ Once the next pixel along the border is located by scanning the neighborhood, the neighborhood area is repositioned and the search for the next border pixel is started. By performing this neighborhood search in the clockwise direction and then selectively shifting the neighborhood, all pixels bordering the object within that slice can be found and their (x,y) coordinates noted.

Once the algorithm locates all object boundaries in a 2-D slice, processing proceeds on subsequent slices until all of the objects in the data volume are detected and their boundaries stored as adjacent (x,y,z) points where the

Fig 2. A method of border following using a moving 3×3 pixel neighborhood and clockwise search. Successively searched neighborhoods are designated 1, 2, 3, and 4. Pixels marked with the dot are the found border points.

z-coordinate is constant on a given slice. The boundary descriptions are then passed to algorithms that decide which adjacent 2-D boundaries are part of the same objects.^{27,40,41} This surface construction step connects adjacent points in adjacent slices into planar segments (polygons) comprising the exterior and/or interior surface of the object. For all but the simplest anatomical structures this is a complicated problem frequently requiring some operator interaction to avoid incorrect joining of slice-adjacent boundaries. This problem renders the 2-D boundary approach inferior in its automatic operation to other direct 3-D methods.

An elaboration of 2-D boundary following uses the full 3-D data to perform a search similar in concept to the 2-D slice searching.³⁰ However, all pixels in the $3 \times 3 \times 3$ voxel neighborhood are considered. Again, the result of the search is a collection of (x,y,z) coordinates fully describing the surface of the object in three-space. The advantage here is that 3-D objects are assembled during the surface decision and traversal process, thus performing the joining of adjacent boundaries with complete 3-D neighborhood information. With appropriately sophisticated rules of joining, the chance of spurious object connection can be minimized.

Another approach to boundary detection processes the entire image (or 3-D volume) to arrive at a calculation of the image gradient at each point. The gradient is proportional to the change in the pixel (or voxel) values within a specified neighborhood. The gradient is high at tissue interfaces and low within areas of uniform tissue and can be used to locate object surfaces. A full consideration of the use of gradient operators is beyond the scope of this discussion and the interested reader is referred to Cappelletti and Rosenfeld.⁴²

Volume extraction within the 2-D slice locates all pixels on the image that are within the object. Of course, this includes the surface voxels in addition to details about the internal structure. Such a method is useful when objects will ultimately be rendered transparent or sliced to show their inner surfaces (such as the location of the bone canal in a bone reconstruction). To carry out volume extraction within the objects, region growing algorithms were developed.¹³ Like surface detection methods, these algorithms start

with a seed point; however, the seed is within the object and the algorithm exhaustively searches adjacent pixels. Unlike surface traversal, which stays on the border between object and background, region growing exhaustively marks/records all adjacent voxels that are in the object until there are no more contiguous object voxels. Like boundary extraction, volume extraction can be extended to the 3-D image volume, thus using full 3-D information in the voxel decision process.

The result of any of the object-determination algorithms is a list of objects, each defined by multiple surface points indicating the location of a voxel center at (x,y,z). In the case of region growing, all points in the object comprise the object data set. In a typical CT skull series, bone surface voxels may comprise over 100,000 coordinate triplets. This is formidable but considerably fewer than the original 15 million voxels present in a 60-slice study. Volume extraction methods produce somewhere between the 10,000 and 15 million points, usually in the hundreds of thousands of points. Because of this, the processing and storage of surface and volume data must rely on sophisticated techniques for storing and coding coordinate data. This is especially true for 3-D programs developed to run on the mini-computers typically found in CT and MRI scanners.⁴³⁻⁴⁵

Rendering. Rendering transforms object representations into meaningful 3-D scenes. The simplest kind of rendering for 3-D object data is the wireframe display.^{27,46} The wireframe is a collection of lines drawn between adjacent points on the object's boundary. For simple scenes (one or two objects) this is sufficient to convey spatial orientation and simple shape information, especially when real-time motion of the displayed object is performed using input from a joystick or trackball. This type of display is possible because of advances in specialized hardware for the display of vector (line-drawing) information. Typically a viewing session on such a device would let the user see the objects, turn them around, move toward or away from them, and possibly separate them. An enhancement to wireframe display is the addition of depth shading to the drawn lines. This wireframe-in-a-fog look is instructive in conveying near-far relationships, but overall, wireframe displays do not produce enough of a 3-D

appearance to be acceptable as the final form of 3-D data.

The most common type of 3-D rendering is called shaded surface display. This method relies on a lighting model for the realistic rendering of well-defined areas of the object surface (surface patches). Surface patches can be defined mathematically as bounded polynomials but are more frequently defined as polygonal patches of a plane represented by three or more coplanar vertices. Polygons can be constructed from object contour lines on adjacent slices or they can simply be the square faces of the voxel cubes. Lighting models can become quite complex as the algorithm designer attempts photorealism. However, the most basic models based on Lambert's Law of Cosines are composed of a light source, an object surface patch with its surface normal vector, and the observer's eye, all with given locations and orientations.⁴⁷ With all of these known, the amount of light that would strike the observer's eye from the object can be calculated. The determination of surface patch orientation, as represented by the surface normal, is crucial to the production of realistic displays. Effort in the development of new algorithms for 3-D display concentrate on the extraction of surface normal data. Beyond simple reflection, we can add diffuse light components to show details in the shadows and highlighting to give a feeling of surface texture and orientation. Figure 3 summarizes the simple reflection model that provides the basis for most 3-D shaded surface displays.

Lighting models presuppose the existence of surface orientation data. It is instructive to consider the difficult problem of going from objects composed of surface points to objects composed of surface patches. As mentioned previously, the noncubic voxels are usually interpolated to form an isotropic (cubic voxel) image volume. By performing all object extraction operations on the isotropic data, the smallest part of the resulting objects possess a uniform surface patch (a square voxel face) and, hence, a limited number of possible surface orientations. By following this cuberille model, the algorithms, given a light source and an observer's location, can precalculate the light intensity of each of the six possible voxel faces.⁴⁴ These six reflection coefficients represent all of the possible light emanating from

Fig 3. Simple lighting model for an object surface patch (triangle) with surface normal N , light source vector L , and reflected light vector R . Elaborations to the regular light model include highlights and diffused light.

the object at each point. Without this cubic simplification, reflection would have to be calculated for each surface patch in each scene. Thus the well-conditioned, cuberille model allows the six trigonometric reflection calculations of Fig 3 to be performed once for each rendered scene. More advanced applications of this fixed-patch orientation scheme can be found in Cline et al.⁴⁸

Another method for defining surface patches uses data from the slice-boundary object extraction algorithm. These data are composed of object boundary loops on adjacent slice levels, and a surface is constructed from them by aligning the loops and drawing lines between closest points. This connection of contours is a complex process requiring sophisticated algorithms to choose the best surface when connecting different numbers and shapes of contours on each slice level. The lines connecting adjacent boundary level points make up the edges of polygonal plane patches with original voxel locations at the vertices.^{40,41,49,50} Typical polygons are triangular and rhomboidal patches. Once this

polygonization is performed, reflection coefficients can be calculated according to the lighting model. This demands trigonometric calculation for each surface patch. This is much more expensive than the cuberille, calculate-once approach, especially because a realistic rendering of a single, curved bony object typically requires anywhere from 5,000 to 500,000 polygons. As a further complication, some 3-D algorithms rely on mathematical representations of the patches as polynomial or spline functions,^{27,51} which further complicates the lighting model and thus increases calculation requirements.

Once the object's reflection can be calculated at each point there still remains the question of how to handle front and back faces, self-obscuring features, and obscuring objects. Add to these questions the representation of shadows and transparency and one can begin to appreciate the complexity and computational intensity of these programs.

One of the simplest (and most time-consuming) methods of rendering a final 3-D image is ray tracing. This method interposes a display window between the observer's eye and the objects in the scene that acts as a projection surface for the objects.^{9,16,52} The window's spatial resolution and extent is determined by the limitations of the display device. In a typical $1,024 \times 1,024$ color display device the display window is composed of 1 million pixels, each of which must be colored by the ray-tracing algorithm. Less computation-intensive approaches reduce this effort by taking advantage of the ultimate display resolution and cubic voxels.⁴⁸

In tracing a light ray from the light source to the surface of an object, it would appear to reflect from the object, possibly through other objects, and proceed through the display window to the observer's eye. Light rays that fail to strike the object or are reflected in another direction never reach the observer's eye. Rather than worry about all light rays from the source, ray tracing turns the problem around and back-traces rays from the observer's eye. Thus, for each pixel on the display, the ray tracing algorithm creates a line from the eye to the display pixel. The algorithm notes any intersections with the scene objects. If all objects are opaque, only the nearest object is considered, and the ray is then traced back to the light source for calculation of the reflected intensity. If the nearest object is transparent, the ray is diminished and possibly refracted to the next nearest object and so on until the light is traced back to the source. In this manner each display window pixel is assigned a light intensity that reflects all objects behind that small display area. Thus an image is built up point by point. Figure 4 shows this process diagrammatically for a two-object scene with transparency.

The amount of calculation necessary can be appreciated by considering a simple example of a scene with two objects, each composed of 10,000 polygons, which is to be displayed on a color $1,024 \times 1,024$ screen. The display window would contain 10^6 points, and for each point we would have to perform 20,000 intersection calculations. For intersecting polygons, we must calculate distance. If we do not have transparency, then we

Fig 4. Method of ray tracing showing how a ray is traced back from the observer's eye through the display window and onto or through an object. Note the addition of a transmission vector T used to account for partial transmission of light through an object.

need only calculate the reflected light intensity three times, one for each color (red, green, and blue) of the colored light source. This results in $10^6 \times 20,000$ intersection calculations plus approximately $2 \times 20,000$ distance calculations, followed by approximately $3 \times 10,000$ reflection calculations. The addition of transparency, multiple light sources, and shadows increases these calculations by orders of magnitude. It is no wonder that commercial animation studios rely on Cray supercomputers for the generation of their photorealistic 3-D graphic images.

Volumetric Imaging

The most recent development in 3-D image reconstruction is the broad application of volumetric imaging,^{9,53} a voxel-based method rooted in the early reprojection (MPR) work of Harris et al.³⁴ Originally commercialized by PIXAR, Inc (San Rafael, CA), this method has been picked up and modified by other manufacturers of 3-D display stations and CT scanners.

This method combines interpolation, tissue mapping, gradient calculations, and a simple lighting model to produce a specialized reprojection of CT data in which soft tissue is seen as a colored haze surrounding translucent bone. In volumetric imaging, the operator first reviews the CT slice data and classifies tissue according to probabilistic intensity maps. Unlike the strict threshold technique used in object-based methods, volumetric rendering allows a voxel to be classified as partially tissue type A and partially tissue type B. Additionally, each tissue type is given a designated opacity, usually with bone

most opaque (95%) and soft tissue least opaque (5%). The tissue classification tables are used to decide the total amount of opacity to assign to a pixel as it is reprojected for display. While simple voxel opacities are being calculated and rays traced through them, a parallel process calculates the surface characteristics of the object using a gradient operator (which estimates the surface normal) and a light source to provide a simple surface. This gradient image is then combined with the opacity image to produce a final rendering in which displayed pixel intensity is due to the accumulated tissue types and surface proximity of all image voxels along the ray path. Figure 5A shows two cranial defect cases rendered with the volumetric image algorithm.

CREATING ANATOMICAL MODELS

One of the most exciting applications of 3-D imaging is the creation of life-size physical models of patient anatomy. Model fabrication brings the industrial techniques of computer-aided-design (CAD) and computer-aided-manufacturing (CAM) to medical data. For the past 15 years, patient prosthetic devices have been designed with the aid of computers.^{54,55} Generally this design process has relied on the manual input of anatomical data, usually from a digitizing tablet. In the past 5 years CT data have been used to directly manufacture models of patient anatomy using numerically controlled (NC) machine tools. Because of the complexity of human anatomy and the limitations of the machining process, this still requires some manual CAD

Fig 5. Two volumetric reconstructions of the skull made from 2-mm interval CT head scans. (A) A patient with a repaired cranial defect and blunt trauma damage to the orbital rim, and (B) a patient with a right temporal defect prior to cranial implant surgery.

interaction. However, sophisticated hardware and software are increasingly automating the process. The next 5 years should see the introduction of 3-D model fabrication systems that, once loaded with a CT data set, can produce bone models without any operator intervention.

The most promising area of 3-D model making is in the field of stereolithography. This technology arose out of the need to build rapid part prototypes for the automotive and aerospace industry. Created by 3D Systems, Inc (Valencia, CA), the system uses an ultraviolet laser to create an object layer by layer by "painting" the surface of a liquid polymer. Every point that is exposed to the layer will solidify. Thus an object representation in a computer is transformed into a life-size object in the model space in the vat of liquid. This is especially appropriate for CT slice data that have been processed into the binary object/no object representation.⁵⁶

DISCUSSION

The use of 3-D imaging in medicine will increase as computer hardware and storage costs decrease and interconnectability of image modalities increases. Although it is not yet clear which 3-D imaging algorithms are best for which cases, there are some conclusions to be drawn by looking at the past 10 years of activity. More and more scanners will use software on the CT and MRI consoles to produce occasional, cost-effective 3-D reconstructions.³⁸ 3-D imaging is likely to migrate away from the acquisition units (CT, MRI consoles) as independent workstations become less expensive and direct digital interfaces to scanners improve. These workstations will integrate specialized hardware for processing, display, and measurement,⁵⁷ accommodat-

ing most of the basic processing and display technologies mentioned in this report. 3-D image output will be in the form of static monochrome images, static color images, dynamic (videotape) color images, stereoscopic displays,¹⁰ holographic images,^{58,59} and anatomical models. The imaging of bone and the simulation of surgery on it will become increasingly sophisticated.⁶⁰ Soft tissue 3-D imaging will become more prevalent as methods of integrating modalities such as CT and MRI become possible. MRI protocols and processing will be developed to allow improved 3-D soft-tissue display and model construction.

Advances in CAD/CAM will allow data to be quickly and efficiently transformed into models of patient anatomy, with precision limited by the imaging modality, not the fabrication device. Model making will increase in sophistication as new technologies allow for the full representation of internal structure, not just surface structure. This will allow the surgeon to actually perform surgical procedures on a full-sized model of patient anatomy. Cutting into the plastic "bone" will show the true anatomy including canals, foramen, and sinuses.

Overall, the prospect for the surgeon, referring physician, and, most importantly, the patient are quite good. Costs can be controlled by deciding the proper level of 3-D service, ranging from the on-scanner software to dedicated processing and display hardware and ultimately to the model fabrication hardware. The radiologist remains as the premier interpreter of the primary images but the nonradiologist can now develop an understanding of the physical relationships of anatomical structures. This increase in understanding should improve surgical planning and communication about patient treatment.

REFERENCES

1. Bloch P, Udupa JK: Application of computerized tomography to radiation therapy and surgical planning. *Proc IEEE* 71:351-355, 1983
2. Fuchs H: An introduction to Pixel-planes and other VLSI-intensive graphics systems, in Earnshaw RA (ed): *Theoretical Foundations of Computer Graphics and CAD*. New York, NY, Springer-Verlag, 1988, pp 675-687
3. Drebin R, Carpenter L, Hanrahan P: Volume rendering. *Comput Graphics* 22:65-74, 1988
4. Goldwasser SM, Reynolds RA, Talton DA, et al: Techniques for the rapid display and manipulation of 3-D biomedical data. *Comput Med Imaging Graphics* 12:1-24, 1988
5. Herman GT, Liu HK: Display of three-dimensional information in computed tomography. *J Comput Assist Tomogr* 1:155-160, 1977
6. Coin CG, Herman GT: The use of three-dimensional computer display in the study of disc disease. *J Comput Assist Tomogr* 4:564-567, 1980
7. Pate D, Resnick D, Andre M, et al: Perspective: three-dimensional imaging of the musculoskeletal system. *AJR* 147:545-551, 1986
8. Fishman EK, Magid D, Ney DR, et al: Three-dimensional imaging and display of musculoskeletal anatomy. *J Comput Assist Tomogr* 12:465-467, 1988
9. Levoy M: Direct visualization of surfaces from com-

puted tomography data. *Med Imaging II*, SPIE 914:828-839, 1988

10. Henri C, Collins L, Peters TM, et al: Three-dimensional interactive display of medical images for stereotactic neurosurgery planning. *Med Imaging III: Image Processing*, SPIE 1092:67-74, 1989

11. Kennedy DN, Nelson AC: Three-dimensional display from cross-sectional tomographic images: An application to magnetic resonance imaging. *IEEE Trans Med Imaging* 6:134-140, 1987

12. Levin DN, Hu X, Tan KK, et al: Surface of the brain: Three-dimensional MR images created with volume rendering. *Radiology* 171:277-280, 1989

13. Stern RL, Cline HE, Johnson GA, et al: Three-dimensional imaging of the thoracic cavity. *Invest Radiol* 24:282-288, 1989

14. Altman NR, Altman DH, Wolfe SA, et al: Three-dimensional CT reformation in children. *AJR* 146:1261-1267, 1986

15. Cline HE, Dumoulin CL, Hart HR Jr, et al: 3D reconstruction of the brain from magnetic resonance images using a connectivity algorithm. *MRI* 5:345-352, 1987

16. Keller PJ, Drayer BP, Fram EK, et al: MR angiography with two-dimensional acquisition and three-dimensional display. *Radiology* 173:527-532, 1989

17. Vandermeulen D, Smets C, Suetens P, et al: Knowledge-based 3-D segmentation of blood vessels on a spatial sequence of MRI and ultrasound images. *Med Imaging III: Image Processing*, SPIE 1092:142-152, 1989

18. Watanabe Y, Nose Y, Sanefuji S, et al: A method for volume estimation by two-dimensional echocardiography: Examination with excised animal left ventricles. *Comput Biol Med* 15:61-69, 1985

19. Fessler JA, Macovski A: Three-dimensional reconstruction of vessels with stenoses and aneurysms from dual biplane angiograms. *Med Imaging III: Image Processing*, SPIE 1092:22-32, 1989

20. Kitamura K, Tobis J, Sklansky J: Estimating the 3-D skeletons and transverse areas of coronary arteries from biplane angiograms. *IEEE Trans Med Imaging* 7:173-187, 1988

21. Pekelsky JR: Automated contour ordering in moiré topograms. *Optical Engineering* 26:479-486, 1987

22. Potmesil M: Generation of 3D surface descriptions from images of pattern-illuminated objects. *Proc IEEE Pattern Recognition Image Processing PRIP-79:553-559*, 1979

23. Addleman D, Addleman L: Rapid 3D digitizing. *Comput Graphics World* November: 41-44, 1985

24. Levin DN, Pelizzari CA, Chen GTY, et al: Retrospective geometric correlation of MR, CT, and PET images. *Radiology* 169:817-823, 1988

25. Peters TM, Clark JA, Olivier A, et al: Integrated stereotactic imaging with CT, MR imaging and digital subtraction angiography. *Radiology* 161:821-826, 1987

26. Chuang K-S: Comparison of interpolation methods in 3-D image display. *Med Imaging IV: Image Processing*, SPIE 1233 1990 (in press)

27. Vandermeulen D, Suetens P, Gybels J, et al: Surface reconstruction from planar cross sections for 3-D integrated

representation of medical images. *Advance Image Processing*, SPIE 804:183-190, 1987

28. Heyers V, Dengler J, Meinzer H-P: Approximation of missing sections of CT-image sequences using binary interpolation, in Magnenat-Thalmann N, Thalmann D (eds): *New Trends in Computer Graphics*. New York, NY, Springer-Verlag, 1988, pp 537-545

29. Ney DR, Fishman EK, Magid D, et al: Interactive real-time multiplanar CT imaging. *Radiology* 170:275-276, 1989

30. Artzy E, Frieder G, Herman GT: The theory, design, implementation and evaluation of a three-dimensional surface detection algorithm. *Comput Graphics Image Processing* 15:1-24, 1981

31. Fishman EK, Magid D, Robertson DD, et al: Metallic hip implants: CT with multiplanar reconstruction. *Radiology* 160:675-681, 1986

32. Johnson GA, Korobkin M: Image techniques for multiplanar computer tomography. *Radiology* 144:829-834, 1982

33. Glenn W Jr, Johnston R, Morton P, et al: Image generation and display techniques for CT scan data. Thin transverse and reconstructed coronal and sagittal planes. *Invest Radiol* 10:403-416, 1975

34. Harris L, Robb R, Yuen T, et al: Display and visualization of three-dimensional reconstructed anatomic morphology: Experience with the thorax, heart, and coronary vasculature of dogs. *J Comput Assist Tomogr* 3:439-446, 1979

35. Perkins W: Area segmentation of images using edge points. *IEEE Trans Pattern Analysis Machine Intell PAMI-2:8-15*, 1980

36. Smets C, Vandermeulen D, Suetens P, et al: Knowledge-based system for the 3-D reconstruction and representation of the cerebral blood vessels from a pair of stereoscopic angiograms. *Med Imaging III: Image Processing*, SPIE 1092:130-138, 1989

37. Kalvin AD, Peleg S: Three-dimensional multiresolution segmentation algorithm for surface reconstruction from computerized tomography (CT) data. *Med Imaging III: Image Processing*, SPIE 1092:173-182, 1989

38. Herman GT: Three-dimensional imaging on a CT or MR scanner. *J Comput Assist Tomogr* 12:450-458, 1988

39. Seitz P, Rueggsegger P: Fast contour detection algorithm for high precision quantitative CT. *IEEE Trans Med Imaging MI-2:136-141*, 1983

40. Christiansen HN, Sederberg TW: Conversion of complex contour line definitions into polygonal element mosaics. *Comput Graphics* 12:187-192, 1978

41. Fuchs H, Kedem ZM, Useton SP: Optimal surface reconstruction from planar contours. *Graphics Image Processing* 20:693-702, 1977

42. Cappelletti JD, Rosenfeld A: Three-dimensional boundary following. *Comput Vision, Graphics, Image Processing* 48:80-92, 1989

43. Herman GT: Three-dimensional imaging in medicine. *Photomethods* 29:24-27, 1986

44. Udupa JK: Display of 3D information in discrete 3D scenes produced by computerized tomography. *Proc IEEE* 71:420-431, 1983

45. Chen L, Herman GT, Reynolds RA, et al: Surface shading in the cuberille environment. *Comput Graphics Appl* 5:33-43, 1985
46. Vannier MW, Marsh JL, Warren JO: Three dimensional CT reconstruction images for craniofacial surgical planning and evaluation. *Radiology* 150:179-184, 1984
47. Foley JD, van Dam A: *Fundamentals of Interactive Computer Graphics*. Reading, MA, Addison-Wesley, 1983
48. Cline H, Lorensen W, Ludke S, et al: Two algorithms for the three-dimensional reconstruction of tomograms. *Med Phys* 15:320-327, 1988
49. Stimac GK, Sundsten JW, Prothero JS, et al: Three-dimensional contour surfacing of the skull, face, and brain from CT and MR images and from anatomic sections. *AJR* 151:807-810, 1988
50. Cline HE, Lorensen WE, Ludke S, et al: High resolution 3D reconstruction of tomograms. *Radiology* 161: 243, 1986
51. Browne MA, Gaydecki PA: High-speed spline fitting, with application to boundary tracing in low-contrast digital images. *Comput Biol Med* 17:109-116, 1986
52. Kuchkuda R: An introduction to ray tracing, in Earnshaw RA (ed): *Theoretical Foundations of Computer Graphics and CAD*. New York, NY, Springer-Verlag, 1988, pp 1039-1060
53. Levoy M: Display of surfaces from volume data. *IEEE Comput Graphics Appl* 8:29-37, 1988
54. Kilgus D, Finerman G, Kabo M: Three dimensional imaging techniques for the reconstruction of a complex hip problem. A case report. *Orthopedics* 11:1029-1034, 1988
55. Rothman SL, Glenn W, Rhodes M, et al: Individualized prosthesis production from routine CT data. *Radiology* 157:177, 1985
56. Mankovich NJ, Cheeseman AM, Stoker NG: The display of 3-D anatomy with stereolithographic models. *Med Imaging IV: PACS Systems and Evaluation*, SPIE 1232 1990 (in press)
57. Robb RA, Heffernan PB, Camp JJ, et al: A workstation for multi-dimensional display and analysis of biomedical images. *Comput Methods Programs Biomed* 25:169-184, 1987
58. Budinger TF: 3D image displays bring scans to life. *Diagnostic Imaging* March:54-56, 1983
59. Boudreaux JC, Lettieri TR: Toward real-time animation of holographic video images. *Visual Commun Image Processing II*, SPIE 845:462-468, 1987
60. Udupa JK: Simulation of surgical procedures using computer graphics. *Proc Natl Comput Graphics Assoc* 3:80-92, 1986