REALISTIC PRESENTATION OF THREE-DIMENSIONAL MEDICAL DATASETS

R. A. Reynolds, P. Decuypere,
S. M. Goldwasser, D. A. Talton and E. S. Walsh

Dynamic Digital Displays, Inc.
3508 Market Street
Philadelphia, PA 19104, USA

B. A. Holshouser and E. L. Christiansen

Loma Linda University Medical Center
Loma Linda, CA 92354, USA

ABSTRACT

Existing strategies for visualizing three-dimensional (3-D) medical datasets such as computed tomography (CT), magnetic resonance (MR), and positron emission tomography (PET) are reviewed and evaluated. Many shaded surface display algorithms are based on segmentation techniques whereby voxels (volume elements) whose densities (grey-scale values) fall outside preset thresholds are simply discarded. To achieve data reduction and thus gain speed, the thresholding step is usually applied early in the processing sequence. However, as this paper demonstrates, premature thresholding introduces errors, including loss of spatial resolution, that prevent faithful rendition of fine surface detail. An alternative approach, based on first-level ray tracing with object re-sampling by trilinear interpolation, locates the bone/tissue interface with sub-voxel precision. Postprocessing with a gradient shading operator results in exceptional 3-D realism when the ray tracing approach is employed.

KEYWORDS: Computed tomography, three-dimensional imaging, volume rendering, ray tracing, gradient shading.

INTRODUCTION

Many medical imaging techniques — such as computed tomography (CT), magnetic resonance (MR), and positron emission tomography (PET) — are inherently three-dimensional in nature. Shape and compositional information about the interior of the human body is collected as a sequence of two-dimensional (2-D) cross-sectional slices making up a three-dimensional (3-D) volume of interest. By mentally stacking these slices into their correct spatial locations, the trained physician can deduce the 3-D structure of individual organs and infer their inter-relationships in space.

A number of specialized computer graphics techniques have evolved to assist this mental reconstruction over the last ten to fifteen years. Display of the original 2-D slices is often augmented by *multi-planar reconstruction (MPR),* a method of creating new slices at arbitrary orientations by interpolation from the 3-D dataset [1, 2]. A related approach creates new images by *reprojection:* data points interpolated along projector lines have their intensities summed to give the "look-through" or shadow appearance of conventional x-rays [3, 4]. However, the most realistic way of presenting 3-D information is by creating *shaded 3-D views* of selected objects such as bony structures, isolated from their surroundings [5, 6]. Shaded graphics presentations — the major topic of this paper — are particularily useful in evaluating major fractures [7] and skeletal joint anomalies [8], and in surgical [9] and radiation therapy planning [10].

REVIEW OF 3-D DISPLAY METHODS

There are 3 main classes of algorithms for generating 3-D views of medical datasets:

1. *Contours* can be traced (either manually or automatically) on individual 2-D slices; these are then rendered as line drawings or wire-frame outlines, using standard vector graphics techniques.

2. *Surfaces* enclosing objects of interest can be formed from contours by triangulation, or directly from the 3-D dataset using automatic surface detection algorithms; these are then displayed using standard shaded graphics techniques such as surface rendering.

3. *Volumes* can be displayed directly using "volume rendering" techniques which do not rely on either contour or surface formation. Instead they operate directly on the original voxel dataset as acquired from the CT scanner or other medical imaging system.

These methods are categorized by the underlying object representations because the representation drives the rendering algorithm, and also determines the flexibility for user interaction.

Reducing the data and discarding some information — as in contour or surface formation — may lead to faster image generation, but also makes the discarded information unavailable or inaccessible for additional processing. Note also that the further the object representation is removed from the raw data format, the more pre-processing is needed (which represents a substantial time delay). The voxel representation retains the full spatial and contrast resolution of the original data, and volume rendering methods avail of this by providing maximum flexibility and faithful rendition of fine detail. Multiple presentation formats — such as

MPR, reprojection, and shaded surface display — can all be generated with volume rendering techniques, and cut-away viewing is facilitated since internal as well as surface information is retained [11].

Despite the growing interest in medical applications of 3-D display, most algorithms have not reached their full potential in terms of image quality and faithful rendition of fine detail. Many widely used algorithms [12, 13] operate by replacing the grey-scale (12-bits per voxel) CT density with a binary (1-bit per voxel) representation of objects or tissue types of interest (typically bone). This segmentation step is incorporated early in the processing for computational expediency (because of the data reduction achieved). However, the binary or thresholded representation results in a loss of spatial information, replacing smooth outlines with jagged silhouettes [Figure 1]. If thresholding is applied prematurely, this jagged or blocky appearance is carried through to the final 3-D views.

The 3-D display algorithm presented in this paper differs from traditional approaches in several respects. First, it does not render 3-D objects from binary or surface representations — instead it retains the grey-scale CT numbers and therefore avoids the errors introduced by premature thresholding. Second, the grey-scale (12-bit) dataset is re-sampled at display time to obtain sub-voxel spatial information [14, 15], thereby revealing fine surface detail. Third, the method is an "image space" approach based on the notions of ray tracing, in contrast to previous methods such as surface rendering [12], Front-to-Back [13] or Back-to-Front [16] which operate in "object space" [Figure 2]. An important consequence is that suitable data can be presented at high magnification without the appearance of objectionable artifacts.

RAY TRACING APPLIED TO MEDICAL DATASETS

Ray tracing has been used to produce some of the most realistic images in mainstream computer graphics. In medical applications, realistic portrayal of the information in the acquired data — not generating aesthetically pleasing images — is the desired goal. When applied to medical datasets [17, 18, 19] the ray tracing approach results in several advantages that are difficult or impossible to obtain with other methods. These include: (i) preservation of spatial information to the sub-voxel level; (ii) elimination or reduction of "jaggies" around the object silhouette; and (iii) reduction of sampling errors in the form of holes or blocky structures at high magnification. Ray tracing retains the advantages of other volume rendering approaches: (i) slice planes can be introduced to reveal internal structure, and (ii) multiple disconnected object fragments can be displayed in their correct spatial locations without additional processing. This in contrast to surface detection methods, which typically select and display one individual object from the 3-D scene.

Any algorithm to render a 3-D scene by projection onto a 2-D display surface must perform several functions, including the following:

1. a coordinate transformation to translate, rotate and scale the scene to the desired orientation;

2. visibility determination or hidden surface removal — elimination of those object points that should be obscured from the observer's viewpoint;

3. perspective or parallel projection onto the viewing plane;

4. realistic shading to compute the appropriate intensity for each pixel in the final image.

The ray tracing approach provides these functions by emulating the physical behavior of light rays (or x-rays). From the eye of the observer a ray is traced through the center of each image pixel until it encounters the surface of an object of interest, thus accomplishing visibility determination and perspective or parallel projection [Figure 3]. The coordinate transformation is inherent in the orientation of the object with respect to the display screen.

In our implementation, "ray tracing" is a greatly simplified form of the sophisticated procedure used in more traditional areas of computer graphics. This is because the large size of medical datasets — 16 million voxels or more — would make more elaborate methods prohibitively time-consuming. For example, orthographic projection is almost always employed in medical applications — light rays are assumed to be parallel to each other and perpendicular to the screen. This is to reduce the rendering time, and also to preserve distances and angles for direct measurement without distortion. A further simplification is the adoption of first-level ray tracing — each light ray terminates when it encounters a voxel belonging to an object of interest. Higher level ray tracing would follow secondary light rays from the point of intersection to one or more (simulated) light sources, emulating the laws of reflection and refraction at all surfaces encountered.

Our first-level ray tracing algorithm with orthographic projection establishes a line perpendicular to the screen through the center of each pixel. It then steps along this ray, testing the CT number at multiple sample points to see if an object of interest (for example, bone versus soft tissue) has been encountered. If so, the stepping terminates, and a shading computation is invoked to assign an appropriate value to the image pixel.

Of paramount importance is the criterion used to terminate the loop along a ray (e.g., when the ray is judged to have intersected the surface of an object of interest). This requires testing the tissue type at precise intervals along the ray, which in turn means that CT numbers must be estimated (interpolated) *between* the collected sample points (voxels). The particular interpolation method selected has important repercussions on the image quality and fine detail apparent in the 3-D rendition:

1. zeroth order (nearest neighbor) interpolation assumes the CT number is constant within the interior of a given voxel [Figure 4a]. Traditional 3-D display methods that threshold the data during pre-processing and/or form surfaces prior to display implicitly make this assumption. The scene appears blocky (especially if magnified), and no sub-voxel resolution is possible.

2. first order (trilinear) interpolation takes the weighted average of the 8 sample points (voxels) surrounding the point of interest [Figure 4b]. A thresholding function is applied *after* re-sampling to obtain the precise location of the object surface. Figures 6 through 8 show images obtained using trilinear interpolation.

3. higher order interpolation (quadratic or cubic spline; sinc function) is feasible but the limited advances in image quality appear outweighed by the additional computational complexity.

The simplest (diffuse) shading model emulates the illumination of a dull or matte surface due to a single light source located at the observer's position [Figure 5a]. The method we have devised to implement this model [20] is called *gradient shading,* which can be applied as a post-processing step [Figure 5b]. Related shading methods, which must however be applied at rendering time, are described in [15] and [17].

DATA ACQUISITION AND DISPLAY

Data were collected with a GE 9800 CT Scanner [General Electric Medical Systems Division, New Berlin, Wis]. Cross-sectional images of an excised femur were taken from a dry bone scanned in air; temporomandibular joint (TMJ) images were from a fresh frozen cadaver specimen; images showing the spine were from a live patient with a history of osteoporosis. Slice thickness was 1.5 mm with 1 mm slice spacing, with the exception of the spine (3 mm slices with 3 mm slice spacing). No gantry angulation was used. All axial slices were reconstructed using the ReView™ facility (bone algorithm); pixel size varied from 0.27 mm (femur) to 0.49 mm (spine).

Axial (2-D) slices were transferred to a MicroVAX II computer [Digital Equipment Corporation, Maynard, Mass.] via magnetic tape for 3-D reconstruction and display. Each 3-D view was generated from the axial data by computer software using the following three steps (which the software follows automatically). First, a low resolution distance shaded image (depth map) was generated at the desired object orientation, using the Front-to-Back [13] or Back-to-Front [16] rendering algorithm. The low resolution depth map served as a starting-point for the ray tracing procedure, which was carried out next (thus avoiding unnecessary computations). The output of the ray tracing procedure was a high resolution depth map which was used as input to the final procedure (gradient shading).

Desired object orientations were obtained by specifying three rotation angles, about X', Y' and Z' axes fixed with respect to a stationary observer. To selectively display bone versus soft tissue, thresholds in the range 200-300 Hounsfield Units (HU) were applied *after* density resampling. All images were generated on a 512 x 512 matrix. The software can generate either perspective or parallel projections; however, parallel (orthographic) projections were used for all the 3-D views presented here.

ReView™ is a trade mark of General Electric Medical Systems

RESULTS AND DISCUSSION

Initial results [Figures 6 through 8] indicate the excellent sub-voxel resolution achievable with the ray tracing approach. Figure 6 serves to demonstrate that the surface detail is a faithful rendition of the original object, and is not due, for example, to artifacts of the computer process.

Of particular importance is the ability to support high magnification or zoom, provided this is warranted by the quality of the input data. As evidenced by Figure 7, the ray tracing algorithm produces highly magnified images without the blocky appearance of individual voxels becoming evident. High resolution studies of the TMJ may be of value in assessing small irregularities of the joint surface which may in turn relate to patient discomfort or pain.

Not surprisingly, the quality of the 3-D views depends in great measure on the resolution and consistency of the axial data used as input. As expected, closely spaced slices produce the best results. The slight ridges apparent in Figure 8 are a result of the relatively large inter-slice spacing (3 mm) in that dataset. Despite these ridges, the details of bony erosion apparent in Figure 8 are quite remarkable and serve to demonstrate that high resolution 3-D views of live patients can be obtained without special scanning protocols.

SUMMARY AND CONCLUSIONS

The premature introduction of thresholding functions in the 3-D rendering process can lead to loss of spatial precision. Ray tracing algorithms can avoid this problem by thresholding the data *after* resampling, thus obtaining sub-voxel resolution. Ray tracing retains the advantages of other volume rendering methods, including cut-away viewing and the display of multiple disconnected object fragments. Highly magnified views can be created without objectionable sampling artifacts.

ACKNOWLEDGEMENTS

REFERENCES

1. W. V. Glenn, R. J. Johnston, P. E. Morton, and S. J. Dwyer, ''Image generation and display techniques for CT scan data,'' *Investigative Radiology*, vol. 10, no. 5, pp. 403-416, September-October 1975.

2. M. L. Rhodes, W. V. Glenn, and Y. M. Azzawi, ''Extracting oblique planes from serial CT sections,'' *J. Comput. Assist. Tomogr.*, vol. 4, pp. 649-657, 1980.

3. L. D. Harris, R. A. Robb, T. S. Yuen, and E. L. Ritman, ''Display and visualization of three-dimensional reconstructed anatomic morphology: Experience with the thorax, heart and coronary vasculature of dogs,'' *J. Computer Assisted Tomography*, vol. 3, pp. 439-446, 1979.

4. W. W. Scott, E. K. Fishman, and D. Magid, "Acetabular fractures: Optimal imaging," *Radiology*, vol. 165, pp. 537-539, 1987.

5. R. S. Ledley and H. K. Huang, "Three-dimensional reconstruction of consecutive transverse axial images from the ACTA-Scanner," *Proc. 28th ACEMB*, New Orleans, Louisiana, 20-24 September 1975.

6. G. T. Herman and H. K. Liu, "Display of three-dimensional information in computed tomography," *J. Computer Assisted Tomography*, vol. 1, pp. 155-160, 1977.

7. D. L. Burk, D. C. Mears, W. H. Kennedy, L. A. Cooperstein, and D. L. Herbert, "Three-dimensional computed tomography of acetabular fractures," *Radiology*, vol. 155, pp. 183-186, 1985.

8. D. Roberts and E. Christiansen, "Quantified three-dimensional imaging techniques for biomechanical analysis of skeletal joints," *Proc. IEEE Engineering in Medicine and Biology Society 8th Ann. Conf.*, pp. 1079-1083, Dallas - Fort Worth, 7-10 November 1986.

9. M. W. Vannier, J. L. Marsh, and J. O. Warren, "Three-dimensional CT reconstruction images for craniofacial surgical planning and evaluation," *Radiology*, vol. 150, pp. 179-184, 1984.

10. D. L. McShan, A. Silverman, D. M. Lanza, L. E. Reinstein, and A. S. Glicksman, "A computerized three-dimensional treatment planning system utilizing interactive color graphics," *British J. Radiology*, vol. 52, pp. 478-481, 1979.

11. D. A. Talton, S. M. Goldwasser, R. A. Reynolds, and E. S. Walsh, "Volume rendering algorithms for the presentation of 3-D medical data," *NCGA'87 Technical Sessions Proceedings*, vol. III, pp. 119-128, Philadelphia, PA, March 1987.

12. L. S. Chen, G. T. Herman, R. A. Reynolds, and J. K. Udupa, "Surface shading in the cuberille environment," *IEEE Computer Graphics and Applications*, vol. 5, no. 12, pp. 33-43, December 1985.

13. R. A. Reynolds, D. Gordon, and L. S. Chen, "A dynamic screen technique for shaded graphics display of slice-represented objects," *Computer Vision, Graphics, and Image Processing*, vol. 38, no. 3, pp. 275-298, June 1987.

14. D. C. Hemmy and T. R. Lindquist, "Optimizing 3-D imaging techniques to meet clinical requirements," *NCGA'87 Technical Session Proceedings*, vol. III, pp. 69-80, March 1987.

15. W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," *Computer Graphics*, vol. 21, no. 4, pp. 163-169, July 1987.

16. G. Frieder, D. Gordon, and R. A. Reynolds, "Back-to-front display of voxel-based objects," *IEEE Computer Graphics & Applications*, vol. 5, no. 1, pp. 52-60, January 1985.

17. K. H. Hohne and R. Bernstein, "Shading 3D-images from CT using gray-level gradients," *IEEE Trans. Med. Imaging*, vol. MI-5, no. 1, pp. 45-47, March 1986.

18. S. M. Goldwasser, R. A. Reynolds, D. Talton, and E. Walsh, "Real-time display and manipulation of 3-D CT, PET, and NMR data," *Proc. SPIE*, vol. 671, pp. 139-149, April 1986.

19. D. S. Schlusselberg, W. K. Smith, and D. J. Woodward, "Three-dimensional display of medical image volumes," *NCGA'86 Technical Session Proceedings*, vol. III, pp. 114-123, Anaheim, California, May 1986.

20. D. Gordon and R. A. Reynolds, "Image space shading of 3-dimensional objects," *Computer Vision, Graphics, and Image Processing*, vol. 29, pp. 361-376, 1985.
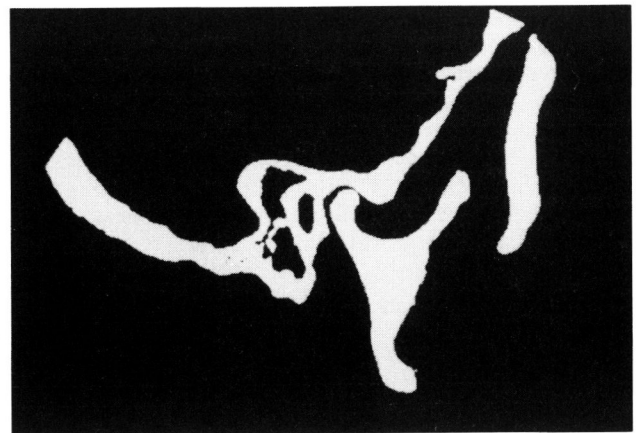
Figure 1 — A single slice through the temporomandibular joint (TMJ), from a direct sagittal CT study of a cadaver head, presented in grey-scale (a) and binary mode (b). Due to the phenomenon known as partial volume effect (PVE), certain voxels at the interface have a value which is an average of soft tissue and bone. These averaged voxels contribute to the smooth bony outline in (a), whereas this "anti-aliasing" effect is lost through thresholding in (b).
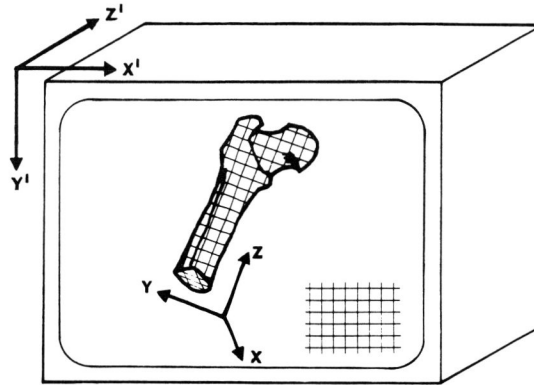
Figure 2 — Object space (X,Y,Z) refers to the coordinate system in which the data are acquired; image space (X',Y',Z') specifies the coordinate system in which objects of interest are presented to the viewer. An object space algorithm loops over voxels making up the object, computing the pixels onto which each voxel projects. An image space algorithm loops over pixels of the image, finding the voxels that project onto them.
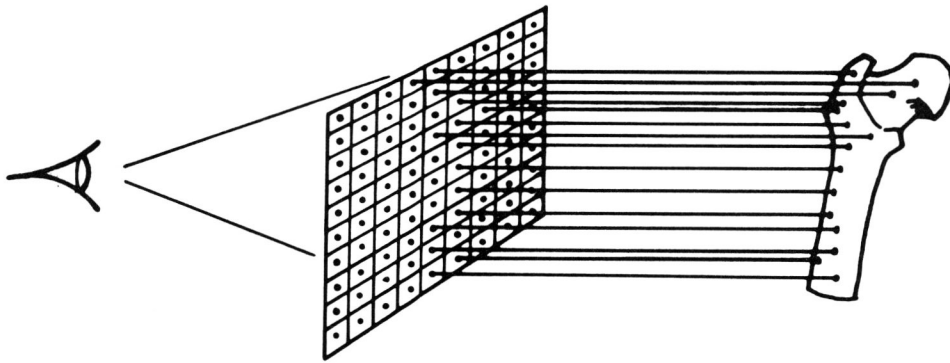

Figure 3 — First-level ray tracing. Light rays are modeled by straight lines traced from the eye of the observer, through the center of each pixel, until the surface of the object of interest is encountered.
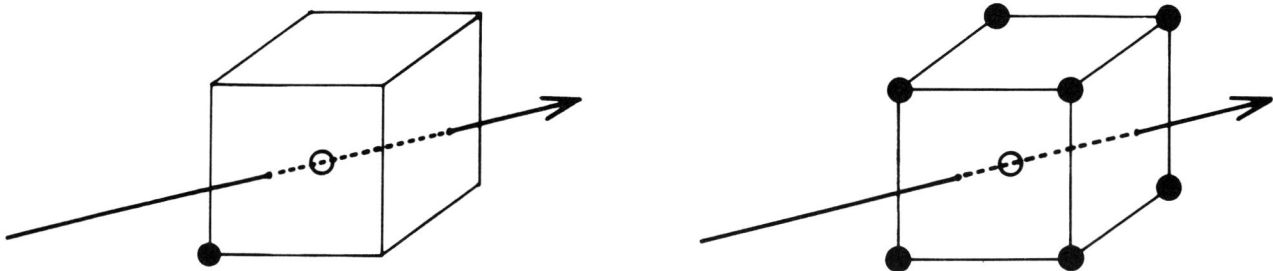
Figure 4 — Re-sampling methods for estimating the CT number between existing sample points (voxels). Zeroth order interpolation (a) uses the value of a single existing sample point; first order interpolation (b) uses the weighted average of 8 existing sample points.
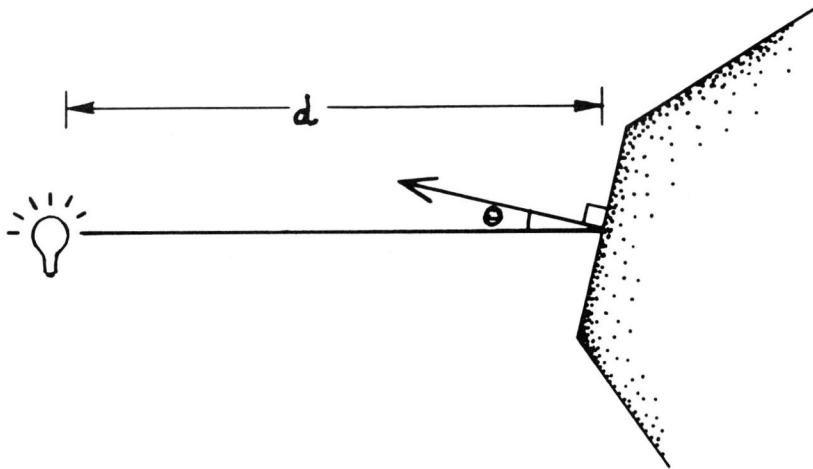
Figure 5a — A simple shading model to emulate diffuse reflection. The surface illumination I depends on the distance d from the light source and on the angle θ which the incident light rays make with respect to the surface normal vector **n**. A simple shading formula is: [20]

$$I = \frac{(I_{max}-I_a)(D-d)(\cos\theta)^p}{D} + I_a \tag{1}$$

where I_a is the ambient light, I_{max} is the maximum intensity supported by the display device, and D is the distance at which the illumination falls to zero (e.g., the diameter of a sphere enclosing the object). p is a parameter that is adjusted empirically (typically 0.2).
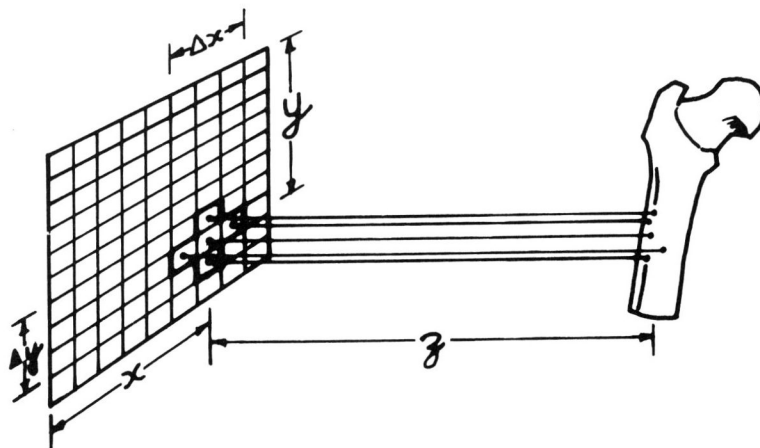

Figure 5b — Gradient shading is implemented with a 3 x 3 operator which estimates the surface normal vectors from a depth map containing the distance $z(x,y)$ to the object surface: [20]

$$\nabla z = \left[\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, -1\right] \tag{2}$$

The advantage of this approach is that it can be implemented entirely as a post-processing step.
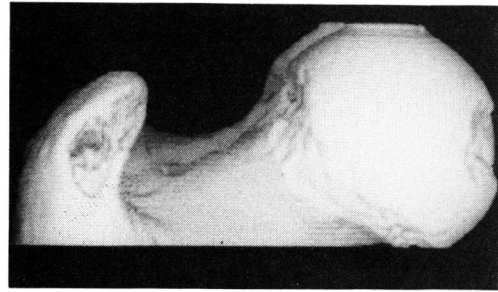
Figure 6 — High resolution computer-generated image of an excised femur (b) compared with a photograph of the same object (a). Note the preservation of surface detail by the ray tracing algorithm.

Figure 7 — High resolution image of a temporomandibular joint (TMJ) generated from CT scans of a cadaver, demonstrating the magnification made possible by the ray tracing approach.
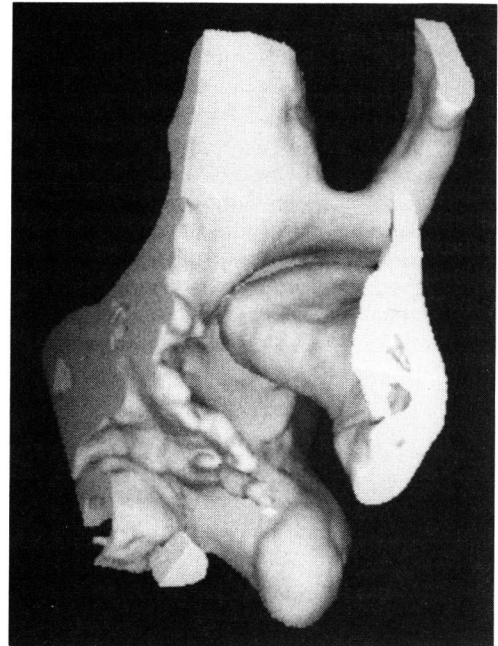

Figure 8 — High resolution image of a single vertebra from a live patient with a history of osteoporosis. Note the detail of bony erosion.