3D SHADED GRAPHICS IN RADIO-THERAPY AND DIAGNOSTIC IMAGING

Technical Report 86-004

S. M. Pizer, H. Fuchs, C. Mosher, L. Lifshits,
G. D. Abram, S. Ramanathan,
B. T. Whitney, J. G. Rosenman,
E. V. Staab, E. L. Chaney, G. W. Sherouse

The University of North Carolina at Chapel Hill
Department of Computer Science
New West Hall 035 A
Chapel Hill, N.C. 27514

3D Shaded Graphics in Radiotherapy and Diagnostic Imaging

Stephen M. Pizer**+, Henry Fuchs*, Charles Mosher*, Lawrence Lifshitz*, Gregory D. Abram*, Sukumar Ramanathan*, Brian T. Whitney*, Julian G. Rosenman +, Edward V. Staab +, Edward L. Chaney +, George Sherouse +

*Department of Computer Science
+ Department of Radiology
University of North Carolina
Chapel Hill, North Carolina 27514

ABSTRACT

Shaded graphics display has been demonstrated to have the potential of providing comprehension of the 3D structure of and interrelationships between organs, diseased anatomy, and surgically implanted objects appearing in medical images such as computed tomograms (CT scans) and magnetic resonance images, as well as of and between calculated objects such as radiation beams and radiation dose distributions. In this paper we present some approaches taken at UNC toward shaded graphics for this purpose, some issues that have arisen, and some results. In particular, we discuss means for object definition, contour association when objects are defined slice by slice, tiling including bifurcation, previewing to set orientation and transparency parameters, high-quality rendering, and presentation.

INTRODUCTION

Producing shaded graphics starting from grey-scale medical images involves the steps of definition of the object(s) to be displayed, tiling its surface, choosing its display parameters, rendering the result, and presenting the rendered image. We will discuss approaches in each of these areas in turn.

OBJECT DEFINITION

In object definition it is traditional to produce object contours slice by slice of the original images. These are frequently hand-drawn, but they may be produced by simple edge-detection techniques such as interactive or adaptive thresholding. Since human interaction is frequently required in this step, convenience and speed are essential properties of the software provided. Our threshold-based autocontouring software (see figure 1) includes not only the means to set thresholds but also means to limit the search for a contour to a specified image region, the ability to delete portions of the contour, and the ability to hand draw portions of the contour and then resume the automatic method to obtain the remainder of the contour.

Following slice-by-slice contouring, decisions must be taken as to which contours are connected to form objects, including where there are bifurcations and end contours. In our system this is done heuristically, based on the centroid of the points on each contour (see figure 2). Editing is then allowed by the user based on a dynamic vector display of the contours and connections in 3D.

Besides slice-by-slice techniques, we are developing methods that operate directly in 3D (or in 2D for single slices and begin with a preprocess that produces an image description of light and dark spots at various scales and with a calculated containment hierarchy [Pizer, 1985; Koenderink, 1984]. The image description is produced by a multiresolution approach that begins by producing successively gaussian blurred versions of the original (3D or 2D) image. It then finds all the local intensity extrema in the image; there will be one spot defined for each extremum. Finally, it assigns

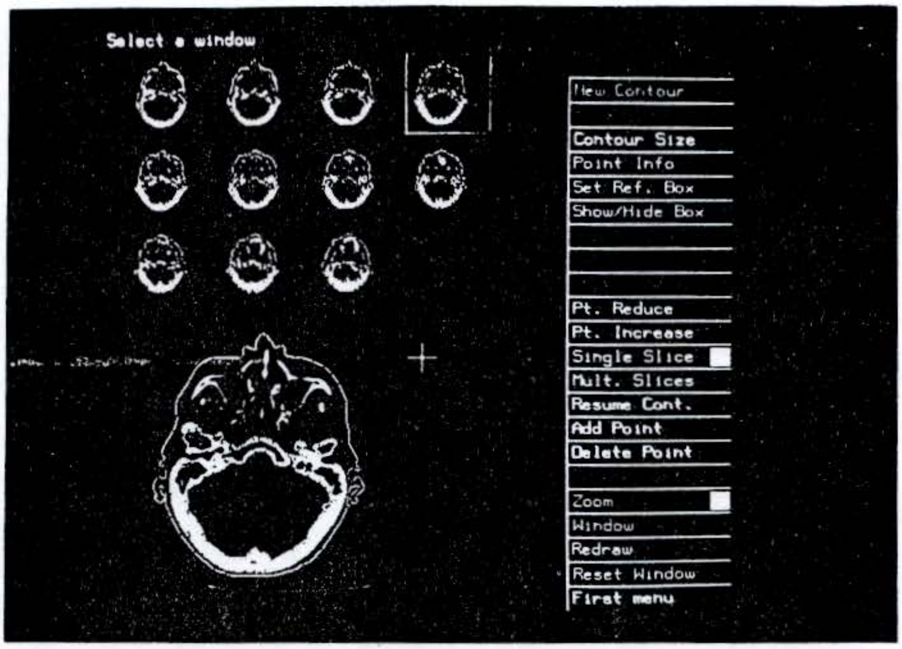

Figure 1. Auto-contouring: a menu and results on CT scan of the Brain.

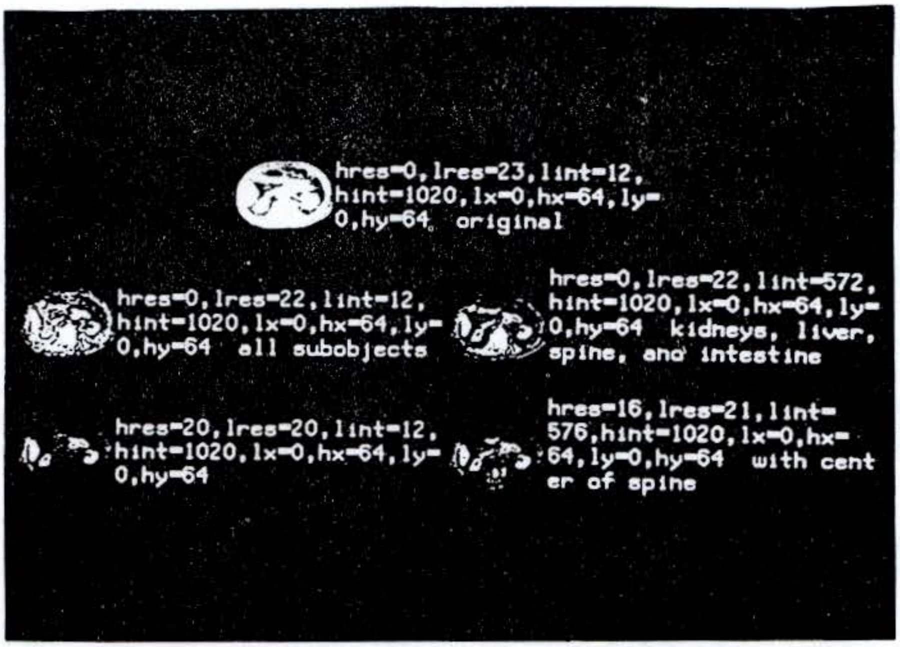

Figure 2. Interactive display based on image description of CT of the upper abdomen. hres and lres represent scale window, lint and hint represent intensity window. Full x and y windows are used.

a scale to each spot as the amount of blurring necessary to make the intensity extremum in the spot disappear into its background and assigns the intensity of the spot as the intensity it achieves when it disappears. This intensity becomes the level of its bounding isointensity contour. Notice that spots include sub-spots that disappeared into them at blurring levels lower than that necessary to make the extremum of the containing spot disappear. All of the above is a preprocess designed to produce an image description tree of spots, each with a scale and intensity as well as a specified group of pixels, and with the tree describing a containment hierarchy.

We are at present using this image description as the basis of a dynamic display component designed to allow the user to determine objects of clinical interest such as organs and lesions. The user specifies the spots making up the object of interest by giving windows in scale, spot intensity, and position co-ordinates. All pixels or voxels that are in spots with values the specified ranges are then displayed at their original intensity, within a few seconds. In a few preliminary studies good, though not perfect, specifications of the regions of interest can be easily determined in a fraction of a minute per object by this means, even for challenging objects made up of soft-tissue in CT scans (see figure 3). He or she can then edit the resulting 2D or 3D figure (set of pixels or voxels). We are also investigating the possibility of model-based pattern recognition by top-down matching of model trees to the computed image description trees.

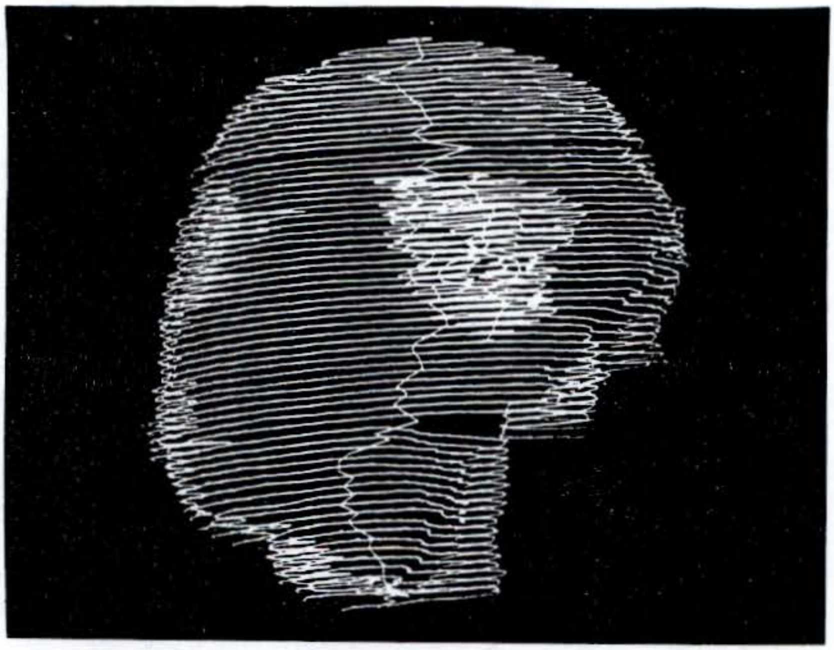

Figure 3. Vector graphics contour connection previewer on contours from CT scan of brain with tumor.

In all of the above steps, as well as the tiling step to be described next, image distances must be measured in world co-ordinates rather than pixel co-ordinates. That is, the relative scales within slices and between slices, the possible offsets of the pixels in one slice from those in another due to nonorthogonality between the slices and the inter-slice offset, and any non-parallelism between slices must be properly handled. The Unix-based image processing package, /usr/image, that we have developed, stores this information about physical dimensions, allowing the world co-ordinates of any pixel to be computed.

A. Tiling a Bifurcation

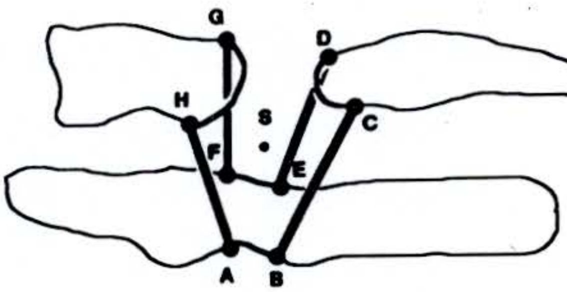

B. Tiling a Bifurcation

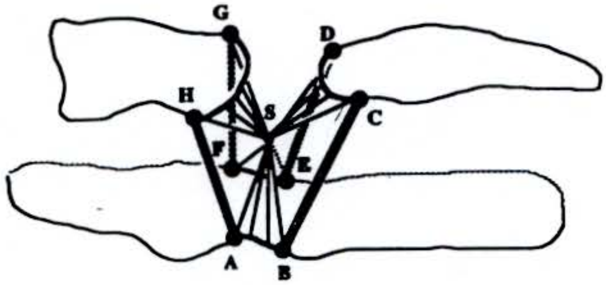

Figure 4. Tiling a bifurcation (a) choosing saddle in spans, (b) tiling. S = saddle point, AB, CD, EF, and GH = spans.

TILING

The process of triangular tiling comes next in our sequence preparing for shaded graphics display. Even if the figure has been defined by a set of voxels, improved display can be obtained if the surface of voxel faces is replaced by a closest fitting skin of triangular tiles. The Fuchs, Kedem, Uselton [1977] tiler has been upgraded by a component that handles bifurcations. In this situation there are two consecutive slices, one with a single contour and the other with two. The approach is to define the saddle region of the bifurcating surface by finding spans on the contours to be tiled that minimize a path length along the boundary of the saddle region (see figure 4). First, a saddle point is defined centrally between the two slices and between the the contours that are paired on one slice. The in-slice coordinates of the saddle point are computed as the midpoint of the line connecting the centroids of the paired contours. The third coordinate is defined to be halfway between the slices. In the limited number of cases where this is unsatisfactory, a new saddle point location can be defined by the user.

Given the saddle point, spans (AB, CD, EF, and GH in figure 4) are found on the three contours to minimize the total distance along these spans and the lines connecting them (ABCDEFGHA in figure 4). In cases where one of the paired contours has no overlap with the single contour on the next slice, this choice of spans may not be satisfactory, and then the user must choose the spans. This choice, and that of the saddle point location, if necessary, are normally made while viewing the contours on a vector graphics display.

Once the spans are specified, tiling the bifurcation is straightforward. Referring to figure 4b, the regions between the spans and the saddle point (ABS, CDS, EFS, and GHS) are tiled with triangles between the successive contour points in the spans and the saddle point. Single triangular tiles (AHS, BCS, DES, and FGS) are defined between the two slices and the saddle. Finally, the Fuchs, Kedem, Uselton tiler is used to tile between the parts of the contours not in the spans (the non-bold path GH to AF and the non-bold path CD to BE).

For a wide range of image objects, the automatic contouring discussed above produces satisfactory results with 10-50 slices in a fraction of an hour and the automatic contour connection and tiling operate without a problem, with the result that a multi-object image can be ready in tiled form in between one and two hours from the time the original medical images arrive. The remaining steps are the choice of orientation to produce a rendered image and the rendering itself.

RENDERING

Rendering a scene made up of a number of tiled objects requires the choice of orientation, as well as color and transparency assignments. This appears to be most effectively done at present using vector graphics, although to obtain a more solid feel while maintaining dynamic control of orientation, viewing distance and angle, and clipping plane, we also use a fast coarse shaded graphics display using ideas of progressive refinement.

We begin by selecting the objects that are to be displayed and computing a binary space-partitioning (BSP) tree [Fuchs, 1983] for the collection of objects. The relative locations of the objects are already determined, because the tiling was done in world co-ordinates. The BSP tree allows the hiding relations to be calculated in time proportional to the number of tiles and a 512 x 512 coarse-shaded image of 250-500 tiles to be calculated and displayed in a fraction of a second. We wish to display this many tiles while interactively choosing the viewing and clipping parameters, so that the display can be reproduced no slower than once every 1/10 seconds. We accomplish this by choosing the integer divisor d of the number of original tiles that leaves 250-500, and dynamically displaying an image from every dth tile. If the interactive devices specifying viewing and clipping parameters have not changed after one display cycle, an image with somewhat more tiles is computed and displayed. This process of progressive refinement continues until either an image from all of the tiles is displayed (see figure 5) or the interactive devices are moved, by an amount that increases over time to make unlikely inadvertent destruction of an image that is time-consuming to recompute.

When the orientation, distance, and clipping parameters, have been chosen by one of the two above-mentioned approaches, these values together with specification of object colors and transparency levels and light source positions are passed to our "high

Figure 5. Dynamic coarse shaded graphics of clipped chest and lungs, from CT scans of chest.

Figure 6. Smooth shaded graphics of skull, brain, and tumor from CT scans.

quality" renderer. This renderer then generates an anti-aliased, Phong-shaded image (from multiple light sources) using a scan-line supersampling approach. By again using the BSP tree priority information, the polygons are rendered in a back-to-front order, enabling it to render transparent objects by adding transparent surfaces into the images as a weighted combination of previous, farther polygons with the nearer transparent surface.

While our present renderer uses the BSP tree as its input, we plan to move to an A-buffer algorithm to handle better the very large number of tiles (hundreds of thousands) that arise in images of clinical complexity.

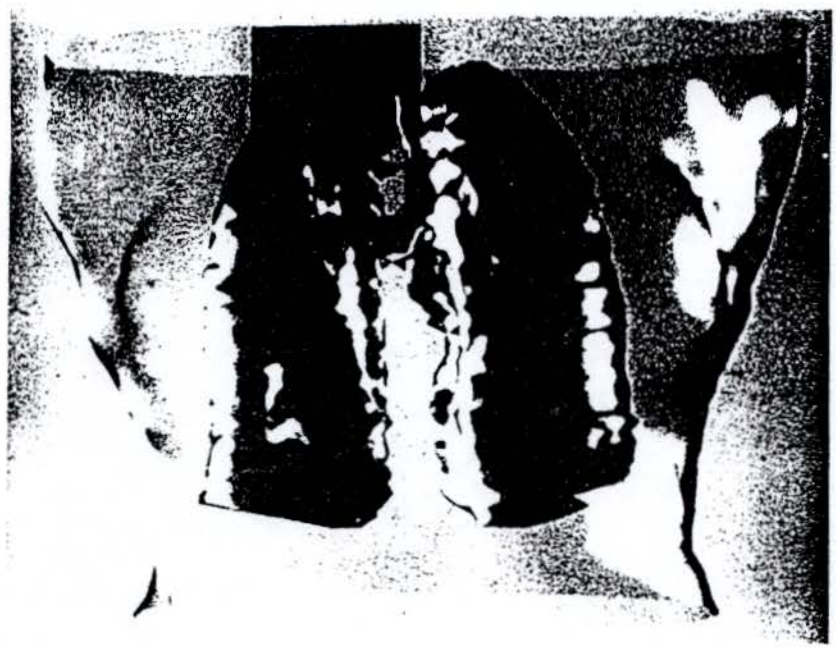

Figure 7. Smooth shaded graphics of radiation treatment beam on chest, lungs, bronchi, and tumor from CT scans.

PRESENTATION

A number of approaches have been investigated to obtain a good 3D percept. We find transparency to be most important, to convey objects enclosed by others (see figures 6 and 7) and the superimposition of computed objects, such as radiotherapy doses or beam profiles (see figure 7) on anatomic objects. Realistic shading with multiple light sources helps to convey three-dimensional information. Smooth shading and good anti-aliasing also appear important for enhancing understanding. Nevertheless, further improvement of the 3D percept is necessary for adequate 3D comprehension. We accomplish this by stereoscopic display or by rotation to take advantage of the kinetic depth effect. While we have found the latter approach to produce very useful results, the computer time required to generate the necessary frames is at present prohibitive (tens of hours). Stereo, on the other hand requires renderings from only two angles, and we find a most effective 3D comprehension to be produced when we view these using the Tektronix stereo plate and presentation of the two images in the odd and even fields of a video display. This plate switches the polarization of the two fields, so the stereoscopic effect can be achieved by viewing with passive, inexpensive glasses whose lenses have opposite polarizations.

RESULTS

We have applied these techniques to a number of CT studies from both diagnostic medical imaging and radiotherapy treatment planning. In the former case these have included not only images of bones, for which the object definition step is relatively straightforward, but in larger numbers images of soft tissues. In particular, we have some experience with images of neoplasms in the brain; organs and vessels in the chest; and the temporo-mandibular joint, including the meniscus. In one case with the chest, the study was ordered by diagnosticians to determine whether a spot was part of a vessel coursing across the slice or part of a tumor, and thus whether an operation on a baby girl was necessary.

In radiotherapy we have been focusing on studies of the brain, pelvis, and chest. The most complex of these to date has included the major bronchi, lungs, and lesions in the chest, as well as the chest surface (see figure 7). Superimposed on these images have been either isodose surfaces or treatment beam profiles.

On the basis of our experience it appears that shaded graphics can be effective in both of these areas of medical imaging. While the work of others [e.g. Herman, 1985] seems to have indicated this fact in bone imaging, our work suggests that it is also true for soft tissue imaging and especially where superposition of anatomy or computed distributions on anatomy are important. Nevertheless, further work is necessary to produce high quality displays that are convenient to produce before this approach will achieve full clinical usefulness.

ACKNOWLEDGEMENTS

We are indebted to Joan Savrock for manuscript preparation and to Bo Strain and Karen Curran for photography.

REFERENCES

Fuchs, H., Z. M. Kedem and S. P. Uselton, "Optimal Surface Reconstruction from Planar Contours", *Communications of the ACM*, Vol. 20, No. 10 (October 1977), pp. 693-702.

Fuchs, H., G. D. Abram, E. D. Grant, "Near Real-Time Shaded Display of Rigid Objects", in *Computer Graphics*, Vol. 17, No. 3 (July 1983), pp. 65-72.

Herman, "Three-Dimensional Computer Graphic Display in Medicine: The MIPG Perspective", *Pictorial Information Systems in Medicine*, K. H. Hoehne, ed., pp. 181-210, Springer-Verlag Berlin, 1985.

Koenderink, J. J., "The Structure of Images", *Biol. Cybernetics, 50:* 363-370, 1984.

Pizer, S. M., J. J. Koenderink, L. M. Lifshitz, L. Helmink, A. D. J Kaasjager, "An Image Description for Object Definition, Based on Extremal Regions in the Stack", to appear in *Inf. Proc. in Med. Imaging,*, 9th IPMI Conf., June 10-14, 1985, Washington, D. C., 1985.