DISPLAY OF MEDICAL OBJECTS AND THEIR INTERACTIVE MANIPULATION

Jayaram K. Udupa
Dewey Odhner
Medical Image Processing Group
Department of Radiology
University of Pennsylvania
418 Service Drive
Philadelphia, PA 19104

Abstract

Three-dimensional imaging is a rapidly developing activity in medical imaging with great potential for improved human health care. The three aspects of this activity – visualization, manipulation, and analysis – together can contribute significanty to both medical diagnosis and therapy, both for routine patient care and for furthering clinical research. Conventional computer graphics techniques have been found not to be appropriate in this activity, as such numerous new methodologies have evolved during the past ten years. Much of the past effort has been on visualization, and to a lesser extent on analysis, and there is very little published work on the interactive manipulation of medical objects and its applications. This paper summarizes the authors' new results in this direction.

KEYWORDS: Medical Imaging, Surgical Planning, Three-Dimensional Imaging, Interactive Manipulation.

I. Introduction

Three-Dimensional(3D) Imaging in Medicine is a multidisciplinary activity that deals with techniques for the visualization, manipulation, and analysis of 3D image data pertaining to human internal organs captured through imaging devices such as computerized tomography (CT) scanners and magnetic resonance imaging (MRI) scanners. The traditional practice in Radiology is to visualize and analyze such data by examining each cross-sectional slice (two-dimensional subset) of the 3D data set either printed on a radiographic film or displayed on the display consoles of the imaging scanners. Visualization and analysis are mostly from a qualitative point of view, and even when quantitative information is sought, the measures are derived from the slices and are hence two-dimensional. With 3D imaging, not only true 3D visualization and analysis have become possible, but manipulation of data – an impossibility in the two-dimensional mode – has opened up new vistas leading to new applications, as we shall try to demonstrate in this paper.

In spite of its appeal, acceptance of 3D imaging in medicine has been rather slow. This is mainly due to the lack of validation studies that establish the superiority of 3D imaging over conventional methods in various clinical applications and due to the high cost of 3D imaging. Two types of 3D imaging equipment are currently available: the CT scanner display consoles, and independent workstations. The former type consists mainly of software, and often hardware, enhancement to the display console that already came with the CT scanner. The price ranges from about $10,000 to over $70,000. They generally have less sophisticated imaging functions and are slower in processing, and hence have less throughput than the workstations. The price and performance of the latter vary over a wide range. They cost from about $50,000 to over $200,000, starting from PC-based systems to specialized display stations with real-time performance. During the past five years there has been a rapid increase in the commercial activity related to 3D imaging as measured from the number of 3D imaging-related products exhibited at the annual meeting of the Radiological Society of North America.

There are mainly two types of computer-display-technology-based approaches [1] to 3D imaging, commonly known as *surface ren-*

dering and *volume rendering.* (There are other lesser-developed approaches based on technologies such as holography and varifocal mirror. These have not yet come up to the level of routine use.) In surface rendering [2-4],object surfaces are explicitly formed prior to creating their depiction on a 2D screen via techniques such as hidden-part removal, shading, translucency, dynamic rotation, stereo projection, and coloring.In volume rendering [5-7], object surfaces are not explicitly computed, rather depictions of surfaces, interfaces, and pseudo-surfaces are generated through a process of cumulative projection of voxels or through ray casting. Most of the published literature in 3D imaging is related to the visualization and analysis aspects. In this article, we shall demonstrate some exciting applications of the manipulation aspect and point out some of the existing problems associated with this aspect.

II. Interactive Manipulation of Anatomy

Visualization, manipulation, and analysis all go hand-in-hand in 3D imaging. While visualization is related to 'seeing' structures, manipulation deals with altering structures, maybe for unobscured visualization or for determining how a deformed structure can be best restored to normalcy. Of course the premise here is that the user directly interacts with the 3D displays of objects on the computer screen. This implies that for interactive processing to be practicable both visualization and manipulation should be doable at least at interactive speeds, preferably in real time. Unfortunately, the requirements for surface display and object manipulation are somewhat contradictory. The main reason is that for surface display we need just the boundary information, since interior points will not contribute to the display. For manipulation on the other hand, we need interior information as well, because when an object is cut, some of its interior points will become boundary points of the resulting object. Our methods are based on an object representation scheme that seems to achieve a good compromise between these requirements.Since a detailed description is beyond the scope of this paper, we give a cursory description of the representation scheme and the methods.

In a preprocessing step we create a representation called a *pseudo boundary* via interpolation, segmentation, and encoding. The pseudo boundary is a set of voxels (*voxel* is an abbreviation for a volume element, a cuboid in general) that belong to the object but that are on its boundary. That is, if a voxel is in the object (as determined by segmentation) and if at least one (but not more than five) of its six face-adjacent neighbors is not in the object then the voxel is in the pseudo boundary. Each voxel in the pseudo boundary has two entities associated with it : (i) a 6-bit binary code that describes which of its six neighbors are in the object, and (ii) a unit vector indicating the normal to the surface of the object at the voxel, computed from the original grey-level image. While the normal is used for shading, the 6-bit code is used in hidden-part removal and all the manipulative operations. For a given view point we can determine from the 6-bit code if a voxel in the pseudo boundary is potentially visible. The actual visibility can be determined either using a back-to-front projection of voxels or using a z-buffer technique. We use the z-buffer technique since the distance map generated in this technique is useful in a variety of operations. The 6-bit code also tells us which side of the voxel the inside of the object lies and which side is the outside. This in fact provides the bridge between boundary and inside (region) information. Since the data structure in which the pseudo boundary is stored allows quick access to the voxels of the pseudo boundary in any given row, we can quickly determine if any given voxel in the row which is not in the pseudo boundary is in the object.In fact this is the key to most of the manipulative operations. For example suppose we wish to cut the object into two parts using a plane. We first compute the 'voxels of intersection' of the plane with each row of the object, and then knowing the voxels of the pseudo boundary and the voxels of intersection compute a new set of voxels representing the pseudo boundary of the modified object. Other complex operations illustrated later on can also be realized using similar techniques.

Whenever possible we carry out a given operation in image space since transforming from image to object space invariably involves the costly scan conversion. For example, when a segment of the object already separated from it through a cutting operation is to be moved (rotated and translated) in the object space, we do so by computing the 3D display of the object and its moved segment and not actually computing the object space pseudo boundary representation of the moved segment. This is sufficient for most purposes except when object space measurements such as volume and intersection of object parts have to be computed.

III. Interactive Surgical Planning

There are many potential biomedical applications for interactive manipulation of discrete 3D objects. The examples we have chosen all come from craniofacial surgery and they illustrate the potential use of interactive manipulation in surgical planning.

Craniofacial deformities are either congenital, or acquired through a disease or an accident. The deformities are usually accompanied by a deformity of the skull as such a basic premise in therapy is to correct the structure of the skull (as well as of the soft tissues) surgically so as to make the face normal esthetically and functionally. Many sources of information are made use of including the surgeon's own experience, life-size photographs, radiographs, CT scan images, and external visual examination. These generally do not give precise 3D information of the extent of deformity nor of the nature and amount of correction required to restore the structure.

With 3D interactive display and manipulation, in principle, any piece of information that is sought after by the surgeon can be derived, though not all of it is possible at the current state of development. In the system that we have developed, we have incorporated a set of tools that we considered are essential in order to evaluate the usefulness of such an interactive system.

Fast display:As we pointed out earlier fast 3D display is vital to interactive processing. In our implementation on a Sun 4/110 workstation with a floating point coprocessor we acheive a 1-4 sec per image speed of display for high-resolution skulls in typical patient data sets. Though this speed is not unsatisfactory, we provide a low-resolution icon derived from the high-resolution object which can be rotated in a fraction of a second as a guide to interactive selection of object orientation (see Figure 1).

Specification of a plane: This can be done for reasons described below either by indicating three non-colinear landmark locations on the surface display (thereby specifying an anatomic plane) or by starting with a geometrically defined plane and interactively positioning it with respect to the object.This also can be done at fast interactive speeds (fraction of a second to compute each new display of the plane with a transparency effect) (see Figure 2).

Mirror reflection: This is an option which may be useful in cases of unilateral deformities. A mirror reflection of the surface about any specified plane can be computed and displayed simultaneously with the original structure in transparent or opaque mode. This operation takes about as long as the display operation (see Figure 3).

Measurement: The distance between any two points on the surface can be measured and so also the angle between the line segments between any three points on the surface, all in real time. The curvilinear distance between any two points on the surface can be measured in a similar fashion by specifying a sequence of points between the two points on the surface.Of course, in the case of mirror reflection, these measurements can be done between the reflected and the original structure.

Cut away: The plane can also be used to cut away parts of the object to reveal its interior (see Figure 4). This operation can be recursively carried out on the modified object. The speed is similar to that of creating the 3D display.

Osteotimes: More complex cutting operations are carried out by specifying the curve of cut on the dsplay and the depth of cut (see Figures 5 and 6). the segment of the object thus separated can be moved freely in image space for repositioning. Once repositioned, the image-space operations such as the measurements and mirror reflection can be carried out on the composite object. The speed of computation of osteotomies depends on the complexity of the curve of cut, but is roughly similar to the speed of display.

Some of the surgical questions require more sophisticated manipulative and analysis tools.One of the questions the surgeon, especially craniofacial, seeks answers for at the time of preoperative planning is what would be the external (soft-tissue) appearance of the patient if a certain procedure were carried out. Another related question is what is the optimum procedure for achieving the best esthetics. A tool that is necessary to answer these questionss (which simply cannot be done using conventional methods) is to drape soft-tissue structures (the muscles, the skin etc.) on the corrected skull structure. Of course, normal functionality of the various organ systems is, if anything, more crucial than esthetics. How the function should be evaluated depends on the nature of the function itself. For mechanical systems this assessment would consist of the analysis of stress, strain and loading, while for other types of systems, for example the visual system, this may have to be further complemented by an optical analysis of some sort.

Acknowledgement

The work reported here is supported by an NIH grant HL28438. The authors are grateful to Ms. Patricia Darby for preparing the manuscript.

References

[1] Udupa, J.K., 3D Imaging in Medicine, Conference Proceedings, Eighth Annual Conference and Exposition of the National Computer Graphics Association, NCGA'87,vol. II pp. 73-104, 1987.
[2] Mazziotta, J.C. and Huang, K.H., THREAD (Three-dimensional reconstruction and display) with biomedical applications in neuron ultrastructure and computerized tomography, American Federation of Information Processing Society, vol.45, pp. 241-250, 1976.
[3] Herman,G.T.,and Udupa,J.K.,Display of 3D information in 3D digital images: computational foundations and medical applications,IEEE Computer Graphics and Applications,vol. 3, pp. 39-46, 1983.
[4] Cook,Lary.T.,Dwyer S.J.III,Batnitzky,S., and Lee,K.R.,A three-dimensional display system for diagnostic imaging applications,IEEE Computer Graphics and Applications,vol.3, pp.13-19 August, 1983.
[5] Harris,L.D.,Robb,R.A.,Yuen,T. and Ritman, E.L., E.L., Non-invasive numberical dissection and display of anatomic structure using computerized x-ray tomography, Proc. SPIE,vol.152, pp. 10-18,1978.
[6] Levoy,M.,Display of surfaces from volume data,IEEE Computer Graphics and Applications,vol. 8, pp. 29-37, 1988.
[7] Drebin, R.A.,Carpenter,L.,Hanrahan,P. Volume rendering,Computer Graphics,vol. 22, pp. 65-74, 1988.
[8] Cutting,C.,Grayson,D.D.S.,Bookstein,F., Fellingham,L.,and McCarthy,J.G.,Computer-aided planning and evaluation of facial and orthognathic surgery, Computers in Plastic Surgery,vol.13, pp. 449-462, 1986.
[9] Udupa,J.K.,Computerized surgical planning:current capabilities and medical needs, SPIE Proceedings,vol.626, pp. 474-482,1986.
[10] Yokoi,S.,Yasuda,T.,Hashimoto,Y., Toriwaki, J.,Fujoka,M.,and Nakajima,H.,A craniofacial surgical planning system, Conference Proceedings, Eighth Annual Conference and Exposition of the National Computer Graphics Association, NCGA, 87, vol.III,pp. 152-161, 1987.

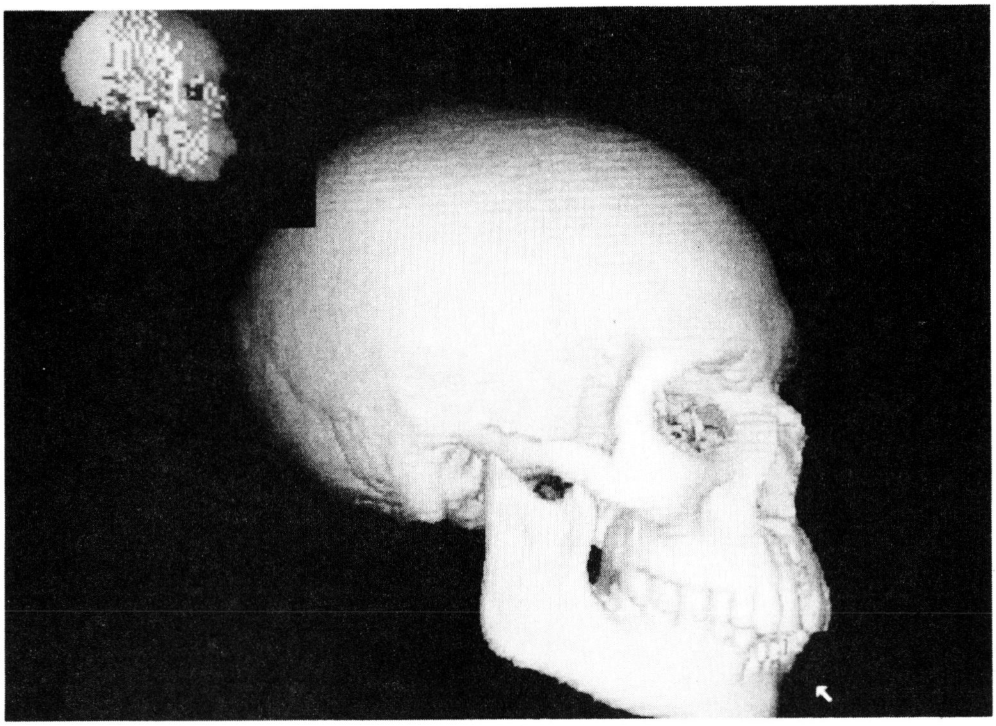

Figure 1. Icon-guided interactive display. The icon (a coarse representation of the actual object) can be rotated interactively to select orientation. We have selected a normal skull to illustrate all operations.

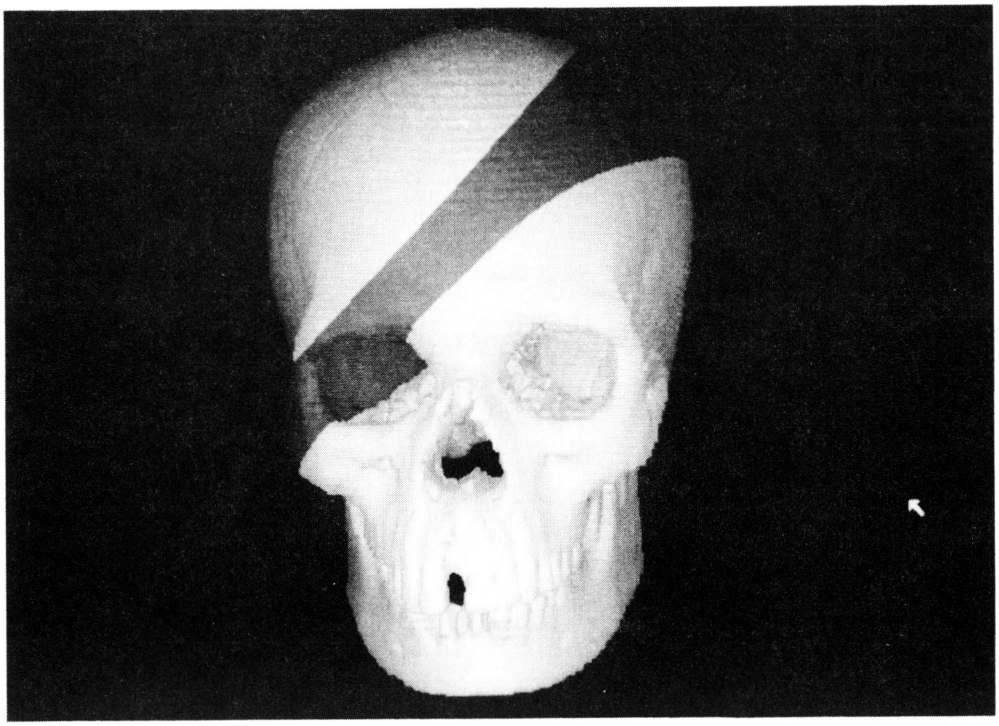

Figure 2. Interactive specification of a plane. A transparent plane can be rotated at interactive speeds keeping the object stationary. In the original color print the plane and the skull are displayed in different colors.

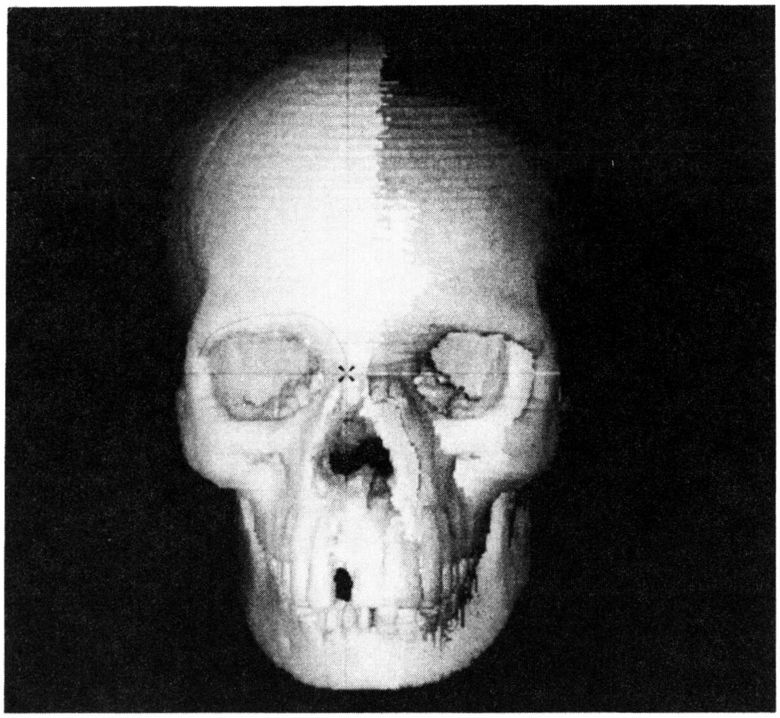

Figure 3. Mirror reflection. The right side of the patient is reflected to the left roughly about a sagittal plane. The display on the patient's right shows both the reflected structure (appearing greyish) and the original structure. Here, both are shown opaque, and hence in some places the reflected structure is visible and in others the original structure is seen (again original color prints reproduced in grey level).

Figure 4. The cut away operation. The part of the skull to the right of the cutting plane shown in figure 2 is retained and shown rotated to the right to reveal its interior.

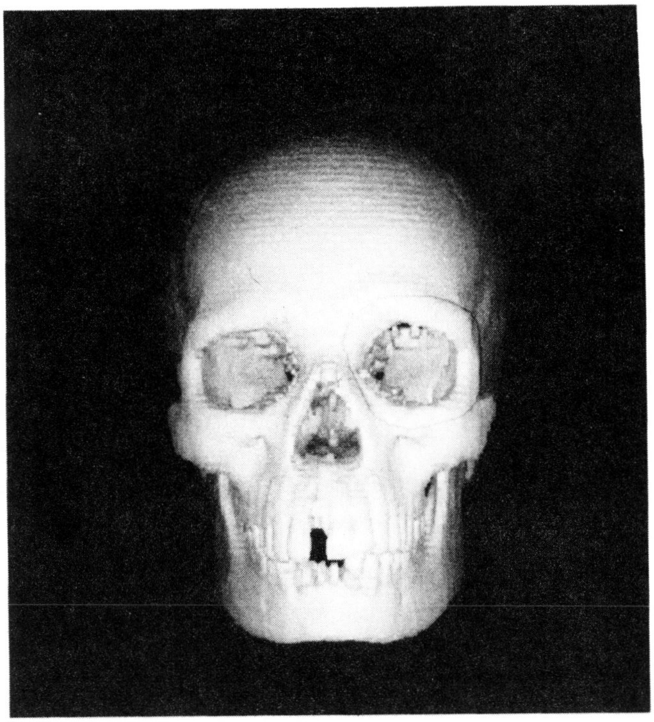

Figure 5. Specification of osteotomy. An osteotomy is specified around the orbit to a certain depth for separating it and subsequent repositioning.

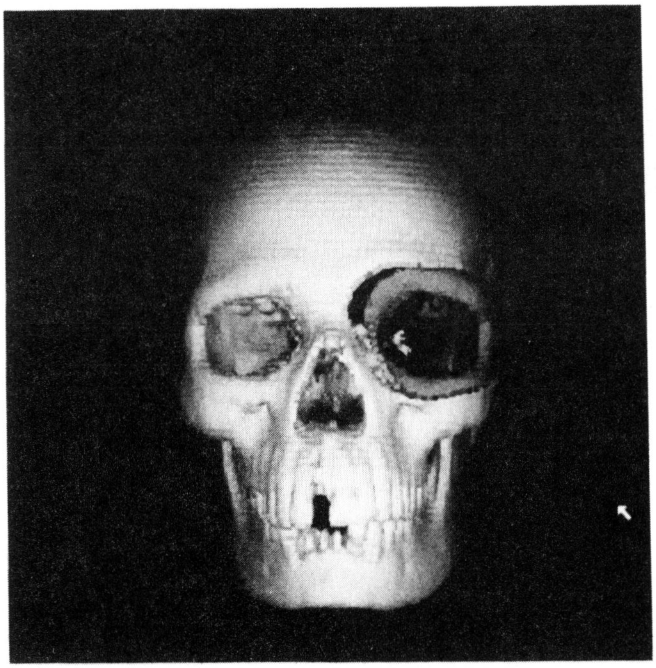

Figure 6. The orbital segment of Figure 5 is repositioned to determine the extent of correction required. In an actual deformity called hypertelorism the orbits are farther apart than normal. The corrective procedure consists of separating the orbits and moving them closer together.