Surface Rendering versus Volume Rendering in Medical Imaging: Techniques and Applications

Organizer
Ramin Shahidi, Stanford University Medical Center

Panelists
Bill Lorensen, General Electric
Ron Kikinis, Harvard Medical School
John Flynn, Silicon Graphics Inc.
Arie Kaufman, SUNY at Stony Brook
Sandy Napel, Stanford University Medical Center

Introduction

High resolution imaging modalities, combined with advances in computer technology has prompted renewed interest and led to significant progress in volumetric reconstruction of medical images. Clinical assessment of this technique and whether it can provide enhanced diagnostic interpretation is currently under investigation by various medical and scientific groups. The purpose of this panel is to evaluate the clinical utility of two major 3D rendering techniques that allow the user to "fly through" and around medical data-sets.

Researchers in the 3D graphics field have been experimenting with new ways to more accurately visualize medical data-sets for the last two decades. As a result, two techniques have emerged which look promising. The first is an advancement in the method of interpreting data-sets by generating a set of polygons that represent the anatomical surface, and displaying a three-dimensional model representation. Polygons representing the outer surface of an object can be calculated using a variant of a "marching cubes" algorithm [Lorensen and Cline]. The method of identifying surfaces of interest, referred to as segmentation, is generally a difficult problem for medical images. The second area of development, volume rendering, is a more direct way for reconstruction of 3D structures. Volume rendering represents 3D objects as a collection of cube-like building blocks called voxels, or volume elements. Each voxel is a sample of the original volume, a 3D pixel on a regular 3D grid or raster. Each voxel has associated with it one or more values quantifying some measured or calculated property of the original object, such as transparency, luminosity, density, flow velocity or metabolic activity. The main advantage of this type of rendering is its ability to preserve the integrity of the original data throughout the visualization process. This technique, however, requires huge amounts of computation time and is generally more expensive than conventional surface rendering technique.

The panelist will use clinical examples to define the future of two methods. Topics to be discussed include the following: which modality will excel for various applications? What are the technical limitations of the techniques and how might they be overcome? Are there significant safety or monetary issues that might favor one or the other technique? Is conventional surface rendering doomed by the need for surface extraction? Will direct volume rendering play only a limited role because of its computational expense? Is there a role for other 3D rendering techniques? and finally what type of navigation techniques ought to be used for each technique?

Techniques for Exploring Large Medical Datasets
Bill Lorensen

The algorithms and hardware technology are here today to produce high resolution surface and volume representations of the anatomy. Two major obstacles are preventing the wide-spread clinical use of 3D: segmentation and user interface. These challenges exist whether we choose a surface or volume representation. Segmentation is the process of identifying organs within the cross-sectional data. This is an active area of research and beyond the scope of this panel.

User interface issues include traditional notions of usability and additional concerns related to working with very large datasets. The issue addressed here is the navigation of large medical datasets. The goal is to provide efficient and usable tools to easily move in and around the narrow tubular passages that exist within representations of anatomy. Three approachs are generally used: 1) Pilot Controls: the user "flys" through the data with a "joystick." 2) Keyframing: the user selects key camera locations with the 3D representation and the system calculates smooth in-betweens for intermediate camera parameters. 3) Path planning: the user provides starting and ending goals and the system automatically finds an "appropriate" path.

The approaches each use techniques from non-medical applications: game technology, computer animation and robot path planning, respectively. If a precalculated path is created by keyframing or path planning, the interface constrains the user to the path but provides interactive controls for stopping and starting. The user can move from the calculated path at any time and interactively explore nearby regions. This is analagous to a guided tour. These navigation techniques are valid independent of surface / volume trade-offs.

Image Guided Therapy: What are the Visualization Requirements?
Ron Kikinis

The trend towards minimally invasive surgery will transform the operating room of the future. In order to minimize the damage to the patient, procedures will be performed increasingly through endoscopic devices. Non-invasive imaging will be required to support the limited localized view from endoscopes. For this purpose, the operating room of the future will be equipped with endoscopes, tracking equipment, intraoperative imaging equipment to update preoperative imaging data, and computer systems that are powerful enough to perform the necessary calculations for registration, tracking, and display of both cross-sectional slices and 3D models.

The presentation will showcase some of the technology involved in image processing, segmentation, visualization, intraoperative registration and tracking. The use of hardware accelerated rendering of surface models will be discussed and demonstrated, using clinical data.

Hardware Technologies Needed for Volumetric Rendering of Medical Data-sets by the Year 2000
John Flynn

The increasing capabilities of computed tomography and magnetic resonance imaging to acquire volumetric data sets with near-isotropic voxels has resulted in increased use of three-dimensional rendering techniques for clinical applications. Three dimensional techniques allow the physician to see anatomical features and interrelationships explicitly and thus make more informed decisions and better communicate them.

Two classes of techniques are commonly used: surface rendering and volume rendering. These techniques although available for years have not been widely used because the cost (compute and/or preprocessing) associated with creating the rendering at interactive speeds have been excessive. However, in the last couple of years a new generation of workstations has brought the benefits of 3D to the clinical community at desktop prices.

As with other computer based technologies, once the constraints on traditional 3D pre-processing and rendering speed are eliminated, new applications using volumetric data in new ways are being developed. One of the most interesting is the virtual endoscopy. The virtual endoscope uses the near iso-tropic data of CT and MRI to view body cavities and hollow structures to be viewed from unique internal perspectives that simulate endoscopy.

One of the limitations of virtual endoscope regardless of technique has been the time needed to prepare or render a fly-though. Widespread use of course depends on the ability to quickly fly-through the data,

Proceedings of the 7th IEEE Visualization Conference (VIS'96)
1070-2385/96 $17.00 © 1996 IEEE

IEEE COMPUTER SOCIETY

even while the patient is on the scanner, and perform the diagnosis. While many are concerned that the compute requirements for clinical use may hold back deployment; these concerns are in fact not warranted.

This panel section includes discussions on: 1) technologies needed for volume and surface rendering , 2) performance levels that are possible today, and 3) the expected performance levels of workstations in the 1998- 2000 time frame. For instance, to volume render a 512x512x150 image set at 10 fps requires a computer capable of rendering 400 million voxels/second. A surface rendering consisting of 2 million polygons requires rendering 20 million polygons/sec to deliver minimal interactivity (10 fps). Larger studies and higher interactivity will require even faster rendering speeds. Graphics supercomputers capable of delivering such performance are available today although at a high price. However; future generation desktop workstations will provide similar performance enabling widespread deployment of virtual endoscope coincident with completion of clinical trials.

Challenges and Developments in 3D Virtual Endoscopy
Arie Kaufman

3D virtual endoscopy has recently been proposed and explored as a non-invasive alternative to the visualization of hollow human organs such as the colon, blood vessel, and intestine, etc. To capture the detailed information from the object of interest, imaging scanners (such as CT or MRI) are commonly used to produce several hundred slices of high-resolution 2D data. This results in relatively large volumetric data and poses new challenges to the visualization process.

With the surface rendering technique, one can convert the volume data into a list of polygons which represents the anatomical surface of interest, using a surface fitting technique such as marching cubes. Due to the high resolution of the volume data, this usually results in an enormous number of polygons. Rendering these polygons with a commonly available graphics workstation can take seconds to minutes. On the other hand, with volume rendering, besides the anatomical surface, one can also visualize the details beneath the surface, which is very helpful especially in differentiating between benign and malignant structures. However, volume rendering even on high-end workstations is still relatively slow. With the advent of specialized volume rendering hardware (e.g., Cube-4 [Pfister & Kaufman 1996]), volume rendering will become, in the near future, the primary mechanism for visualizing medical data.

Besides rendering, camera control is another important issue in 3D virtual endoscopy. Keyframing requires physician's expertise and is tedious and time-consuming, while path planning calculates a "flight" path automatically and possibly generates better viewing angles [Lorensen et al 1995, Hong et al 1995]. However, both keyframing and path planning only provide a general overview of the hollow organ and allow no user interaction once the animation is produced. To examine a local area (which may have been identified with the path-planning animation) in more details from different distances and angles, interactive "walk-through" is essential. Besides interactive rendering speed, how to guide the user through a "journey" of the hollow organ while providing certain flexibilities for the user to move the camera is challenging.

In the last few years, we have been conducting research on 3D virtual endoscopy, primarily focusing on virtual colonoscopy. In addition to an automatic path planning algorithm proposed earlier [Hong et al 1995], we have developed techniques to achieve interactive walk-through within the hollow organ. To obtain fast rendering speed, instead of adopting surface simplification techniques to reduce the number of polygons (which may not be appropriate since the camera is so close to the surface and the image quality should not be sacrificed), we have designed a new visibility algorithm by taking advantage of the fact that usually only a portion of the organ is visible from a certain viewpoint location. Whenever a desirable viewing angle is obtained and further examination of a particular structure is needed, we switch to volume rendering to visualize the information hidden underneath the surface. Furthermore, for the camera control of interactive walk-through, we have developed a physically-based model which employs the concept of potential field to achieve guided navigation inside the hollow organ while avoiding collision with the surface. Our initial studies on several cases (including a pipe phantom, the Visible Human data, and several patients) have demonstrated the effectiveness of this technology.

Volume Rendering and Virtual Endoscopy in Radiology: Is it the Holy Grail?
Sandy Napel

Surface Rendering (SR) and Volume Rendering (VR) are two of many different approaches to 3D visualization of medical (and other) image data. SR is generally faster on equivalent hardware, primarily because the imaged volume is first reduced to a more compact set of pre-extracted surfaces. On the other hand, VR retains all of the volume data during rendering, and all acquired voxels may contribute to each rendered image. Spatial gradients may be computed prior to or during rendering, and these gradients, as well as the relative contributions of each voxel, may be adjusted to produce images of surfaces. That is, volume rendering can produce images of surfaces. Therefore, one might conclude that in the not too distant future, when volume rendering of medical data sets is inexpensive and fast enough, surface rendering will no longer be utilized. However, this simplified argument leaves out one important fact: inherent in surface generation is a segmentation step. Thus, although volume rendering can make images of surfaces, it generates images that include ALL of the surfaces that have similar voxel values. This may make it difficult to examine a specific surface.

Recent developments that incorporate perspective allow the user to adjust his or her viewpoint to within the volume for optimum inspection of a specific surface. Thus, during perspective VR of an iodine-enhanced CT examination of the aorta, for example, it is possible for a viewer to enter the chest, adjust the mapping from density to opacity so as to view the blood vessels of the lungs. Next, he or she can readjust the mapping and enter the trachea to explore the airways. Finally, another adjustment can allow for virtual angioscopy of the aorta and its branches. This type of exploration is currently being developed and evaluated using many types of cross-sectional data (e.g., CT, MR, ultrasound) for many clinical applications.

One also must not forget that other approaches to viewing the details of of volume data are also valuable and may remain so. Curved planar reformations allow exquisite views of vascular morphology along tortuous paths. Thin slab renderings have recently shown utility for detection of small pulmonary nodules. It is most likely that one particular rendering approach will not replace the need for all others. The future will more likely see an integration of many approaches into one manageable toolbox so that the viewer can easily switch from a virtual endoscopic view of a constricted trachea to a thin slab rendering through the constriction, in order to reveal the mass and its involvement in surrounding organs. Volume rendering will be an important component in this toolbox, but not the only tool.

Proceedings of the 7th IEEE Visualization Conference (VIS'96)
1070-2385/96 $17.00 © 1996 IEEE