

Visualizing Volume Data Using Physical Models

David R. Nadeau, Michael J. Bailey

San Diego Supercomputer Center
University of California, San Diego

Abstract

Visualization techniques enable scientists to interactively explore 3D data sets, segmenting and cutting them to reveal inner structure. While powerful, these techniques suffer from one serious flaw – the images they create are displayed on a flat piece of glass or paper. It is not really 3D – it only can be made to appear 3D. In this paper we describe the construction of 3D physical models from volumetric data. Using solid freeform fabrication equipment, these models are built as separate interlocking pieces that express in physical form the segmentation and cutting operations common in display-based visualization.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Computational Geometry and Object Modeling – Hierarchy and geometric transforms; Object hierarchies; Curve, surface, solid, and object representations.

Additional Keywords: scene graphs, volume graphics, volume visualization, physical models.

1. INTRODUCTION

Volume data sets provide a sampled representation of measured or computed quantities that vary across a region of space. Medical scanning technologies, for instance, can produce data sets that depict the inner structure of the body. Computer simulations can produce data describing flow across an airfoil, temperatures in the Pacific Ocean, etc. Data sets such as these exhibit complex inner structure embedded within the volume.

To reveal this inner structure, visualization techniques can segment and cut away portions of the data. Segmentation assigns an opacity value to individual data values within the volume. Data of interest is given high opacity, while less interesting data is made transparent. By setting different opacity levels for different data values, users can interactively disassemble the data set to reveal, say, the skull in medical data of a human head. Adjusting opacity levels further can remove the skull and reveal the brain, the eyes, or other data set components.

Figure 1.1 illustrates these common segmentation and cutting techniques. Here data for the head of the Visible Human Male [5] is segmented to extract data for the skull, brain, optic nerves, and

eyes. In Figure 1.1a, a cutting volume slices open the skull and a transform rotates half of it away to reveal the brain inside. In Figure 1.1b, additional cutting volumes slice the brain and spread the slices apart to reveal the brain's complex inner structure.



Figure 1.1. Segmentation and cutting volumes reveal the inner structure of a complex volume data set.

Segmentation and cutting are obviously powerful ways of exploring a data set. However, computer graphics-based exploration has one serious flaw – the images it creates are displayed on a flat piece of glass or paper. The images are not really 3D, only an illusion of 3D created by clever software.

In contrast, real-world physical models have several visualization advantages:

- They are in the real world. A physical model can be held, felt, and rotated about in a natural way.

- Physical models can be viewed interactively regardless of complexity. The real world has no polygons/second limitations. Similarly, lighting, shadows, and collision-detection are all free.
- Physical model viewing doesn't require expensive graphics hardware. The models can be viewed anywhere, even by people without technical training, such as children in a K-12 classroom.

Display-based visualization has had the advantage because of its ability to *non-destructively* segment, cut, and explore data. This is more difficult to do in the real world. Performing segmentation of a real-world head to cleanly extract the skull and brain would be a very destructive (and gruesome) process.

In this work we leverage the strengths of display-based visualization and real-world physical models. We start by using segmentation and cutting techniques to non-destructively extract data of interest from volumetric data. Then we use the results to manufacture multiple pieces for an interlocking physical model of the segmented data. The user can manipulate these models, in the real world, to gain further insight into the data.

2. BUILDING PHYSICAL MODELS

2.1. Segmenting the Data

This work used the head portion of the Visible Human Male data set from the National Library of Medicine [5]. To manipulate the data, CT and cryosection volumes, as well as a commercially available segmentation volume, were imported into a volume modeling system.

The volume modeling system is based upon a *volume scene graph* [4]: a hierarchical organization of shapes, groups of shapes, and groups of groups that collectively define the content of a scene. The scene graph's nodes are *functions* that take a 3D location in space as an argument. When evaluated at such locations, a node function returns a color value. Most functions return values that may vary over space. Procedural texture functions, for instance, use 3D noise, turbulence, or marble vein algorithms to compute a return value. Volume data set functions use the 3D-location argument to select and interpolate voxel values from a data set read from disk.

Voxelization of a volume scene graph repeatedly evaluates the graph's functions, once for each 3D location in a grid spanning a region of interest. Each evaluation returns an RGB-alpha value that is saved into a voxel in a new discrete volume data set.

To visualize the Visible Human Male head, CT, cryosection, and segmentation volumes of differing resolutions were imported as leaf nodes in a volume scene graph. Each data set was left in its original resolution, without resampling.

For the CT data, classification nodes in the scene graph used the data set's scalar values to select RGB-alpha values that gave the skull a high opacity while dropping the rest of the data set to transparency. For cryosection data, scene graph nodes masked the segmentation volume against the cryosection data to extract data for the brain, optic nerves, and eyes. The CT skull was transformed to register it into the coordinate space of the cryosection data. Finally, the skull and brain data were composited together to create a scene that was voxelized and volume rendered. Figure 2.1a shows the rendered scene and Figure 2.1b the scene graph.




Figure 2.1. CT, cryosection, and segmentation volumes combined by a scene graph to create a scene containing the skull, brain, optic nerves, and eyes.

2.2. Isovoluming for Manufacturability

When drawing a graphics scene on a display, an isosurface is an adequate solution. But for manufacturing the fabrication machine must have some sense of what is meant to be solid (inside) and what is meant to be air (outside). Thus, a surface is not enough. Rather, what must be produced is an *isovolume* [3], or what has also been called an interval volume [6]. This process requires two scalar values to be specified, an S_{min} and an S_{max} . The isovolume will be the solid volume that is between these two isosurfaces.

It is interesting to note that while most of the visualization world is concerned with polygon decimation, this process can improve the smoothness of the model by applying "polygon incrementation" [3]. This process recursively subdivides boundary triangles to create a surface that is a closer approximation to the true trilinear interpolation function within each voxel cell. Such things are reasonable when you care more about model quality than interaction speed.

To enable the physical model of the brain to be taken out of the skull, the skull needed to be split in half. Not long ago, the only way to make this happen would have been to generate the isovolume, import it into a CAD solid modeler, and use constructive solid geometry techniques to split the skull and move the two pieces away from each other. The volume scene graph approach, however, is a much more optimal way to go about this. In data handling theory, there is the concept of higher-order data and lower-order, or derived, data. In this case, the original volume data is the higher-order data and the polyhedral isovolume is the derived lower-order data.

When moving from high order to low order data, some information is, by necessity, removed. In this case, it is the

relationship between the data values among the multitude of voxels that is lost. For example, suppose that there was a certain amount of noise in the original volume data set. As a volume, that noise could be filtered out prior to generating the isovolume. But, if the isovolume had been created from the noisy volume first, then chances are slim that the noise could be filtered out of the isovolume's polygons.

The same can be said for resolution control. While the data is in voxel form it can be filtered, subsampled, or supersampled. The fineness or coarseness of the polyhedral isovolume can be controlled very tightly depending on the resolution of the voxel data that produced it. Isovoluming experiments have shown that adding more polygons through the polygon incrementation process does not produce as smooth a result as adding more polygons by supersampling the original volume.

This work used the scene graph to slice the data set into manufacturable pieces. Figure 1.1b shows these pieces, including the skull cut vertically, and the brain cut horizontally into four slices. Arbitrarily oriented flat or curved cutting surfaces could have been used just as easily. Figure 2.2 shows the scene graph to generate these pieces.




Figure 2.2. The scene graph to cut volume-domain data into manufacturable pieces.

2.3. Manufacturing the Pieces

In 1995, the San Diego Supercomputer Center (SDSC) TeleManufacturing Facility (TMF) project was started as a way to produce physical hardcopy to enhance visualization. This project has resulted in hundreds of visualization models being made for researchers all over the country [1,2,3,7].

The TMF consists of two fabrication machines, both connected to the Internet to provide access and monitoring from virtually anywhere. One machine is a Helisys Laminated Object Manufacturing machine, which makes 3D parts from layers of paper or plastic. The other machine is a Z Corporation Z402 machine, which fabricates from layers of powder. Both machines leave their scrap in place during the build so that overhangs are supported and needn't require any perturbations in the part geometry.

All rapid fabrication machines are driven by the de facto standard STL file format. The exact format of the STL files can be found at [8], but basically it is a list of the 3D triangles that bound the outer skin of the solid. The isovoluming program is set up to produce its list of triangles in the STL format.




Figure 2.3. The Helisys 1015 Laminated Object Manufacturing machine.




Figure 2.4. The Z Corporation Z402 machine.

2.4. Results

Using the scene graph in Figure 2.2, the CT, cryosection, and segmentation volumes of the Visible Human Male were segmented and cut into pieces. The individual scene pieces were voxelized, isovolumed, and expressed as STL files. The skull halves were fabricated on the Helisys LOM machine, while each of the brain slices were fabricated on the Z402 machine. Figures 2.5 and 2.6 show the skull and four brain slices. Images of the cryosection data have been adhered to the slices.



Figure 2.5. The fabricated skull and four brain slices within the skull.



Figure 2.6. The fabricated skull and four brain pieces with cryosection slices adhered.

3. CONCLUSIONS

The volume scene graph approach gives a great deal of flexibility to downstream analyses. In the case of 3D-visualization hardcopy, that flexibility results in better control over the location and orientation of a model's individual pieces. It also gives better control over the smoothness of the resulting isovolume. The physical models produced offer a compelling and natural way to explore the interior of complex volumetric data sets.

Acknowledgements

The volume renderer we used was developed by Jon Genetti at the San Diego Supercomputer Center. The scene graph software and volume renderer are part of a larger suite of scalable volume visualization tools being developed under the NSF-funded National Partnership for Advanced Computational Infrastructure (NPACI). The scalable tools project is led by B. A. Pailthorpe and A. Olson. The telemanufacturing work was funded by the National Science Foundation under grant MIP-9420099.

References

- [1] Michael Bailey, "TeleManufacturing: Rapid Prototyping on the Internet," *IEEE Computer Graphics and Applications*, v. 15, no. 6, November 1995, p. 20-26.
- [2] Michael Bailey, Jack Johnson, and Klaus Schulten, "The Use of Solid Physical Models for the Study of Macromolecular Structure", *Current Opinion in Structural Biology*, v. 8, no. 2, April 1998, p. 202-208.
- [3] Michael Bailey, "Manufacturing Isovolumes", *Proceedings of the International Workshop on Volume Graphics*, 1999, p. 133-146.
- [4] David R. Nadeau, "Volume Scene Graphs," *Proceedings of the IEEE Volume Visualization and Graphics Symposium 2000*.
- [5] National Library of Medicine, "The Visible Human Project." www.nlm.nih.gov/research/visible
- [6] Gregory Nielson, "Tools for Triangulations and Tetrahedralizations and Constructing Functions Defined over Them", from Gregory Nielson, Hans Hagen, and Heinrich Muller, *Scientific Visualization: Overviews, Methodologies, Techniques*, IEEE Computer Society, 1997, p. 429-525.

[7] Kathy A. Svilt, "A Touch of Science", *Discover*, v. 19, no. 6, June 1998, p. 80-84.

[8] The TMF has been renamed the *Center for Visualization Prototypes*. For more information on this project and on physical fabrication in general, see: <http://cvp.sdsc.edu>