

Volume Rendering of Segmented Tubular Objects

Elizabeth Bullitt¹ and Stephen Aylward²

Departments of ¹Surgery, ¹²Radiology, and ²Computer Science, University of North Carolina, Chapel Hill, NC
bullitt@med.unc.edu, aylward@unc.edu

Abstract. This paper describes a new method of combining ray casting with segmented tubular objects, such as blood vessels, for purposes of clinically useful display. The method first projects segmented tubes using a modified z-buffer that additionally records information about the objects projected. A subsequent step selectively volume renders only through the object volumes recorded by the z-buffer. In common with traditional “block” volume rendering the actual image data is shown, of importance when the boundary of a segmented object is uncertain. Unlike traditional “block” volume rendering, the approach permits user manipulation of objects, operates rapidly, and provides depth information even for maximum intensity projection. Although our methods were developed for display of the intracerebral vasculature, the approach is applicable to volume rendering of tubular objects throughout the body.

1 Introduction

This paper describes a new method of combining volume rendering and segmentation that maximizes the strengths of each approach. The goal is to provide a method of three-dimensional (3D) visualization that will be useful to clinicians who need to understand complex vascular anatomy in 3D, who may or may not have access to specialized hardware, and who may be distrustful of automated segmentation methods that sharply define objects subsequently shown only by surface rendering.

From the clinical perspective, an ideal method of visualizing 3D volume data should include:

- 1) The ability to visualize “truth” (the actual image data as seen in high quality images),
- 2) The ability to view objects rapidly and interactively from any point of view,
- 3) The ability to interactively remove obscuring objects and/or to obtain relevant information about selected objects or groups of objects,
- 4) The ability to run the software on any platform/hardware so that surgical planning can be done on office machines or at home.

This paper describes an approach that combines segmentation with volume rendering in order to do all of the above. The basic methodology employs segmentation to define regions that are subsequently selectively volume rendered

using a modified zbuffer. "Fuzzy margins" can be set so that desired amounts of the actual image data are shown regardless of the boundary definitions provided by segmentation. No time is wasted casting rays through non-interesting space, so the implementation operates efficiently even without hardware optimization. Individual objects or groups of objects can be interactively selected, queried, and manipulated. Depth information is available even for maximum intensity projection (MIP), thus adding to the image information available for each image frame.

The methods described here were initially developed for visualization of the intracerebral vasculature as extracted from magnetic resonance angiograms (MRA) and as seen by maximum intensity projection (MIP). However, the approach is applicable to display of tubular objects extracted from any type of 3D image data and from any part of the body.

2. Methods

2.1 Image acquisition and image segmentation

For this study, 12 patients underwent 3D, time of flight MRA in a Siemens 1.5 T Vision unit with a quadrature head coil and with magnetization transfer suppression. Images were acquired in the axial projection over a 7.6 cm volume, using 69 1 mm contiguous sections and an x and y spacing of 0.85 mm. An additional patient was studied by 3D digital subtraction angiography (3D-DSA). Siemens has given us a 3D volume representing this patient's right carotid circulation.

We have previously described and evaluated methods of segmenting vessels from 3D images [1], [2] and of providing directed graph descriptions of these segmented vessels [3], [4], [5]. The approach is applicable to any type of 3D data including MR, computed tomography, 3D-DSA, ultrasound, and confocal microscopy. In brief, each vessel must be provided an initial seed point. The program then automatically defines an image intensity ridge representing the vessel's 3D skeleton curve, with a subsequent automatic calculation of vessel width at each skeleton point. Graph description is performed largely automatically, employing both distance and image intensity measurements. We have tested the resultant vessel trees against digital subtraction angiographic data obtained from the same patient [3], [5], [6]. The output of the segmentation/tree definition program is a set of vessels, each of which is comprised of a set of 3D skeleton points with an associated width at each point, and with additional information available about the connectivity of each vessel.

2.2 Visualization

The visualization method consists of two steps: front projection of segmented objects using a modified z-buffer algorithm, and subsequent volume rendering so that a pixel's ray is processed only if that ray intersects a volume of interest. If more than one object projects to the same pixel, volume rendering is only performed for the

object closest to the ray source. This approach provides an efficient means of volume rendering that does not waste time in processing empty space or in analyzing distant objects that project to the same pixel as a closer object.

For tubular objects such as blood vessels, the approach is simple. For segmented vessels, each vessel skeleton point is projected upon the viewplane. If more than one point projects to the same pixel, the point closest to the ray source is given priority.

Tubular objects are circular in cross section. Each 3D vessel point can therefore be viewed as a sphere whose projection will be circular (or close to circular) on the projection plane. Given the depth (Z) coordinate of each vessel skeleton point, knowledge of that point's radius, and knowledge of the viewing geometry, the radius of each projected point can be calculated on the viewplane. A fast Bresenham filled circle algorithm is used to calculate the associated viewplane pixels from which to back project rays. These rays are selectively cast only through the appropriate 3D object of interest, and only for the distance indicated by that object's radius. If another object has already been identified as projecting to a portion of a filled viewplane circle, the object closest to the ray source is given priority.

A byproduct of the approach is that each viewplane pixel of interest is automatically labeled with the identification number of the object projecting to it and with that object's relevant point identification number. Any subsequent "point and click" operation to interactively select an object can thus be processed rapidly without requiring a 3D search.

2.3 Evaluation

We have processed 12 patient MRA data sets and an additional patient studied by 3D digital subtraction angiography (3D-DSA) using the methods outlined above. Two patients were normal, two had base-of-the-brain tumors, one an aneurysm, one an intracerebral hemorrhage, and 7 an arteriovenous malformation (AVM). The intracerebral vasculature normally consists of three vascular trees: two carotid circulations and one vertebrobasilar circulation. All three intracerebral vascular trees were defined for all 12 patients who underwent MRA. Only one vessel tree was available for the patient who was studied by 3D-DSA.

Images were projected using perspective, MIP projection, a field of view of 20 degrees, and a large window of 500 x 500 pixels. The program allows interactive manipulation of the point of view as well as the ability to interactively select objects. All tests and evaluations were done on an 850 MHz, Pentium III laptop computer running Windows 2000.

Image quality was assessed by side-to-side comparison with the excellent quality, MIP volume renderings of the publically available Visualization Toolkit (VTK) [7]. For purposes of comparison, the viewing geometry was set identically in each window, and the volume of space given to VTK for volume rendering was the bounding box of the segmented vessels shown in the other window.

Volume rendering over large volumes, under perspective projection, and using large display windows is usually very slow without specialized hardware. For each of the 12 patients studied by MR (and for whom the full intracranial volume was thus

available for analysis), we tested the time to render the vasculature 10 times into a 500 x 500 pixel display window.

3. Results

3.1 Image quality

MRAs are noisy, and the quality of an image produced by volume rendering is often degraded by noise. An advantage of selective volume rendering through segmented objects is that little or no background noise is present. In addition, traditional methods of volume rendering may suffer from image obscuration when objects of interest and extraneous objects occupy the same block of space. Selective volume rendering does not have this limitation. Figure 1 shows a patient with an intracerebral hemorrhage. The vasculature adjacent to the hemorrhage cannot be visualized by traditional volume rendering, but is readily appreciated following vessel segmentation.

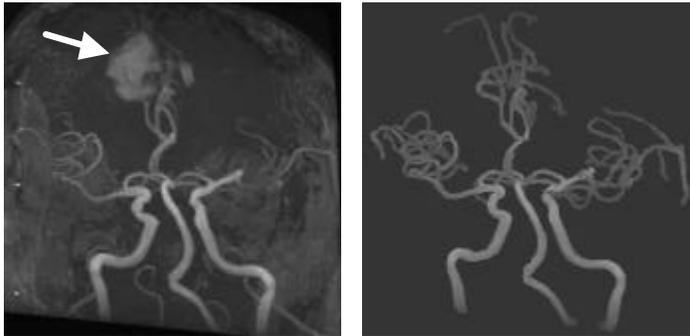


Fig. 1. Comparison of block volume rendering (left) with volume rendering using segmentation (right) in a patient with an intracerebral blood clot (arrow). The image at left is noisy and the pattern of the vasculature is obscured in the vicinity of the clot.

Another advantage of the proposed approach is that depth information is available on any individual view. This leads to images of higher quality and informational content. Figure 2 illustrates the depth information available by our approach. The patient shown is normal.



Fig. 2. Depth information from an individual MIP view. No depth information is available with traditional “block” volume rendering (left), and the resultant image is confusing. Useful depth information is available by the approach outlined in this report (right).

3.2 Image manipulation

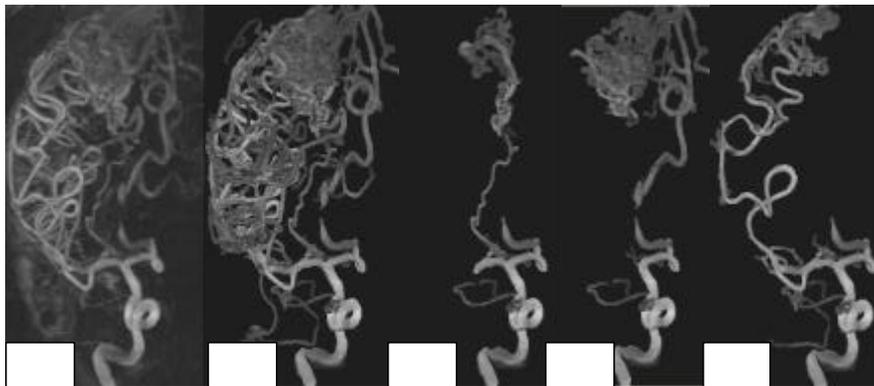


Fig. 3. Block volume rendering (A) and segmented vessel volume rendering (B) of a patient with an AVM defined by 3D-DSA. Our graph descriptions of segmented vessels permit hiding an entire vessel subtree with a single mouseclick. CE illustrate how the arterial feeding patterns can be clarified when subtrees are selectively hidden.

An advantage of segmented object display over traditional block volume rendering is that individual objects or groups of objects can be selected and manipulated by the

user. Figure 3 shows a patient with an AVM whose vessels were extracted from a 3D-DSA. The number of vessels and the degree of projection overlap make it difficult to discern the feeding arteries. By using graph descriptions of segmented objects, a single mouseclick can hide a vessel subtree, thus allowing easier analysis of the vascular feeding pattern.

A disadvantage of segmentation is that the definition of object boundaries may be uncertain or inaccurate. An advantage of combining volume rendering with segmentation is that the actual image data is shown in the region of interest. If the clinician is concerned that segmentation may be inaccurate, our approach allows arbitrary dilation of segmented boundary margins so that more of the actual image data can be seen. Figure 4 shows a patient with a mycotic aneurysm of the middle cerebral artery. Our segmentation did not show an aneurysm neck, and the question was raised whether a neck might actually be present. Figure 4 shows a selective volume rendering of the relevant vessels using a radius of 1.0, 2.0, and 2.5 times the length of each point's radius as determined by segmentation. No aneurysm neck is present, and the initial segmentation was correct. However, the fact that a clinician can check the segmentation against the actual image data is reassuring to the clinician and adds a significant safety factor if 3D volume visualizations are to be used for surgical planning.



Fig. 4. Volume rendering using fuzzy segmentation boundaries. The patient has an aneurysm (arrow) and the question was whether an aneurysm neck might be present. From left to right, volume rendering is shown for 1.0, 2.0, and 2.5 times the radius at each vessel point.

3.3 Image rendering speed

Traditional methods of high quality volume rendering are slow because each pixel on the viewplane must define a ray between itself and the ray source, walk through a volume block, query each voxel it encounters for its greyscale value, interpolate the greyscale values of adjacent voxels, and then process each interpolated value. Under

conditions of traditional “block” volume rendering, much time is wasted in processing voxels without useful informational content. The software algorithm described in this report is capable of fast, interactive volume rendering because it spends time casting rays only through structures of interest. Our approach allows interactive manipulation of object position/orientation on a laptop computer without specialized hardware, using a large display window of 500 x 500 pixels, under conditions of perspective projection, and displaying the entire intracranial vasculature.

Time tests for the 12 images analyzed under the difficult conditions listed above indicated a mean rendering time of better than 5 frames a second. More specifically, the mean time to render an image was slightly less than 0.2 seconds with a standard deviation of 0.1 seconds.

4. Discussion

This report describes a new method of visualizing segmented tubular objects using a modified zbuffer and selective ray-casting. The resultant visualizations allow direct display of image intensity values, the ability to manipulate/query objects interactively, high quality images, fast interaction time, depth information on individual frames, and the capacity to provide “fuzzy” segmented object boundaries so that the actual image data can be shown beyond segmented object margins if the accuracy of segmentation is in doubt.

Other investigators have also approached the issue of selective volume rendering using segmented objects. Like our own method, such approaches offer significant advantages over traditional “block” volume rendering. However, unlike the “shell” rendering approach of Grevera [8], our method does not require explicit construction of object surfaces, allows ray casting precisely through the confines of the objects of interest, and provides high-quality, interpolated images at even high magnification. Unlike the 3D spatial subdivision and labeling approach of Zuiderveld [9], our method avoids walking rays through empty space, allows arbitrary dilation of segmentation boundaries, and permits user-defined manipulation of connected groups of objects. We therefore believe that our approach provides a new and powerful visualization method likely to be useful to clinicians.

A disadvantage of the proposed methodology is that it is applicable only to tubular objects. We are currently exploring methods of extending the approach to segmented objects of any arbitrary configuration.

In addition to the vascular visualizations enabled by the proposed approach, we can also offer significant, hardware-independent increases in rendering speed. Image size appears to be increasing in parallel with hardware improvements. A more efficient software algorithm can take advantage of hardware optimizations and, regardless of the hardware available, should perform better than a less efficient approach. Moreover, a number of algorithms are available that speed orthographic projection. Our current program is implemented only for the more complex case of perspective projection; further speed improvements should be possible under orthographic

visualizations. It should be noted, however, that the rendering speed improvements offered by our approach will be most marked when small objects are scattered over a wide volume of space. If the objects of interest consist only of closely apposed cubic blocks, our approach will operate no more quickly than traditional volume rendering methods, although the other advantages of our methodology will still apply.

Acknowledgments

Supported by R01CA67812 NIH-NCI, P01CA47982 NIH-NCI and Intel and Microsoft equipment awards. Algorithms for vessel segmentation and description have been licensed to Medtronic Corp. (Minn., Minn) and to R2 Technology (Los Altos, CA).

References

1. Aylward, S.R., Pizer, S.M., Bullitt, E., Eberly, D.: Intensity ridge and widths for 3D object segmentation and description IEEE WMMBIA IEEE 96TB100056 (1996) 131-138.
2. Aylward S.R., Bullitt E.: A comparison of methods for tubular object centerline extraction. (2001) Accepted IEEE-TMI.
3. Bullitt E., Aylward S., Liu A., Stone J., Mukherji S., Coffey C., Gerig G., Pizer S.M.: 3D graph description of the intracerebral vasculature from segmented MRA and tests of accuracy by comparison with x-ray angiograms. IPMI 99; Lecture Notes in Computer Science (1999) 1613:308-321.
4. Bullitt E., Aylward S., Bernard E., Gerig G.: Special Article. Computer-assisted visualization of arteriovenous malformations on the home pc (2001) Neurosurgery 48: 576-583.
5. Bullitt E., Aylward S., Smith K., Mukherji S., Jiroutek M., Muller K.: Symbolic Description of Intracerebral Vessels Segmented from MRA and Evaluation by Comparison with X-Ray Angiograms (2001) in press Medical Image Analysis.
6. Bullitt E., Liu A., Aylward S., Coffey C., Stone J., Mukherji S., Muller K., S, Pizer S.M.: Registration of 3D cerebral vessels with 2D digital angiograms: Clinical evaluation Academic Radiology (1999) 6:539-546.
7. Schroeder W., Martin K., Lorensen B.: The Visualization Toolkit. Prentice Hall, New Jersey (1998). Open source software is available at <http://www.kitware.com/vtk.html>.
8. Grevera G., Udupa J., Odhner D.: An order of magnitude faster isosurface rendering in software on a PC than using dedicated, general purpose rendering hardware. IEEE Transactions of Visualization and Computer Graphics (2000) 6:335-345.
9. Zuiderveld K.: Visualization of multimodality medical volume data using object-oriented methods. Thesis. Universiteit Utrecht, Utrecht. ISBN 90-393-0687-7 (1995).