

Direct Volume Rendering from Photographic Data

David Ebert¹, Tim McClanahan², Penny Rheingans¹, and Terry Yoo³

¹ Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County
[ebert,rheingan]@cs.umbc.edu

² Laboratory for Extraterrestrial Physics
NASA Goddard Space Flight Center
xrtpm@lepxgrs.gsfc.nasa.gov

³ Office of High Performance Computing and Communication
National Library of Medicine
yoo@nlm.nih.gov

Abstract. Direct volume rendering from photographic volume data has the potential to create realistic images of internal volume structure, as well as the structure of boundaries within the volume. While possession of the photographic volume simplifies color calculations in voxel illumination, it complicates opacity calculation. This paper describes a framework for addressing illumination challenges in photographic volume data and presents initial results.

1 Introduction

In recent years, a few photographic volume data sets have become available. The most widely used of these are those of the Visible Human Project (VHP) at the National Library of Medicine[15], but other examples are being created as well. This type of data offers exciting possibilities for realistic volume visualization, since correct color values are known for each voxel. Applications include medical illustration, surgical simulation, and general scientific education. Photographic volume data also offers a challenge to traditional volume rendering techniques.

Traditional direct volume rendering produces an image from a volume of scalar data, using transfer functions from scalar value to color and opacity. For some common data types, such as medical CT volumes, designing an opacity transfer function can be straight-forward. One common approach is to assign high opacity to voxels containing high scalar intensity values, for instance those containing bone. The exact transfer function from density to opacity can be modified to suit the purpose of a particular visualization, but one with monotonically increasing opacity is generally effective. For other data types, such as medical MRI volumes, effective opacity transfer functions are more difficult to design, but the basic problem remains one of mapping one scalar value into another.

The design of the color transfer function, from the scalar value to an RGB triple, is more open-ended. One transfer function might try to map each density value to the color of the material of that density in order to produce a realistic image. If each density uniquely corresponded to a single material type, and each material was a uniform

color, such a transfer function would be well-defined. Unfortunately, in medical volumes, both conditions are violated, making realistic color recovery from density values problematic.

Realistically visualizing photographic volume data, where each voxel has a vector RGB value, turns the traditional transfer function design problem inside-out. Now, the color transfer function becomes trivial, since the actual color at each voxel is explicitly known. Design of an effective opacity transfer function is now the major challenge, since neither individual nor average color component values have a natural correspondence with desired opacity.

One fairly successful approach for rendering from photographic volume data has been to compute surfaces or opacity values from an auxiliary volume, for instance a CT volume, and use the photographic data simply for color information. The use of surface techniques results in good representation of the outer boundaries of the object, but reveals little or nothing of the internal structure. With either auxiliary volume approach, the density volume must be registered with the photographic volume, a challenging process which is difficult to automate.

We describe a framework for volume rendering directly from photographic data without the need for a secondary scalar volume to indicate density. Direct volume rendering from the photographic volume eliminates the need to register volumes from different sources and enables the display of internal volume structure along with material boundary information. We describe results of our preliminary experiments using this framework and suggest fruitful directions of continued inquiry.

2 Photographic Volume Data

Photographic color data is becoming increasingly important as a volume information representation. While the acquisition of such data often requires the loss of physical integrity of the sample, tomographic sectioning has been important in anatomy and in pathology. With the growing capability of aggregating multiple photographic cutplane images into 3D volumes, these techniques are growing in importance in other fields. Limitations of MR imaging to generate adequate resolution led the Whole Frog Project at Lawrence Berkeley Laboratory to create an entire frog dataset using cryosection [13]. The use of this dataset is growing as a basis for teaching dissection in biology. The Laboratory for Neurological Imaging at UCLA commonly uses cryosectioning to gain the resolution and contrast required for their intricate analysis of the brain's pathways [16]. Commercial groups now offer mechanical sectioning technology that is used in volume data analysis for geology, medicine, and semiconductor manufacturing [5]. The National Library of Medicine produced one of the most important examples of data acquired through photographic tomographic imaging.

The Visible Human Project was formed to explore the use of digital imaging technology in modern anatomy research and education. A panel study in 1990 recommended as a first project the construction of a digital image library of volumetric data of a normal adult human female and male subject. Data from two subjects, one male and one female, were collected through a variety of methods including the conventional radiological techniques of X-ray CT studies, magnetic resonance imaging MRI, and plain

film radiographs. In addition to these conventional clinical studies, the subjects were frozen and sectioned at 1 mm (male subject) and 1/3 mm (female subject) intervals. The exposed surfaces were photographed with 35 mm and 70 mm film and digitized with an electronic camera. Image acquisition was carefully performed and the resulting data is one of the most complete anatomical studies ever performed [15].

Each slice of the digital cryosection data was acquired with a raster resolution of 2048 x 1216 pixels with a horizontal field of view of 25 inches. Voxel dimensions are 0.32 x 0.32 x 1mm in the male dataset. Reduced cost of image storage, networking, and data handling combined with the desire for cubic voxel dimensions prompted more aggressive data acquisition of 0.33 x 0.33 x 0.33 mm voxel resolution in the female dataset. Each subject was encased in frozen gelatin, dyed blue with food additives, to provide physical stability as well as inter-slice opacity. As each slice of the frozen specimen was exposed, the surface was cleaned and sprayed with ethanol to reduce diffuse reflections from frost, masked to eliminate glare from the insulating materials, and a color platen was placed in the field of view for reference. The resulting surface was then photographed with two film and one CCD cameras. The block was dressed with a tray of dry-ice, to refreeze the exposed surface, re-sectioned, and the process repeated. Voids exposed through the sectioning process were filled with blue latex to stabilize the walls, prevent the collection of debris, and to block projection of deeper structures onto the image plane. The blue tint was chosen to ease the subtraction of non-anatomical structures added during the data collection process. Figure 1 shows a sample cryosection cross-section, transecting the eyes and optic nerves of the male subject. The voids of the nasal cavities have been filled with blue latex. Embedded fiducials are visible in the corners. A platen with CYM color dots and an intensity gradient are visible at the bottom of the image for color and grayscale correction. Over 1800 visible light images were acquired for the male, and over 5000 visible light images were collected for the female subject. The resulting datasets are 14 gigabytes and 40 gigabytes respectively.

3 Related Work

Volume rendering as a image generation and reconstruction technique was pioneered in computer graphics by Levoy [9] as well as by Drebin et al. [2]. Both of these early papers included medical data acquired using X-ray CT as examples. X-ray attenuation is a physical property that is particularly amenable to assigning a color transfer function. Levoy in particular explicitly includes color assignment and shading in his rendering pipeline. These initial implementations and much of the derivative work used pseudocolor to improve the visualizations produced through volume rendering. A comprehensive survey of volume rendering is provided by Kaufman [7]. More recent work employing hardware 3D texture graphics capabilities for volume rendering more easily accommodates color data [1]. However, these methods must still generate an opacity transfer function to enable correct compositing of the image planes. Working from segmented data simplifies the problems of opacity assignment, but includes a huge labor investment in hand segmentation. Recent renderings of segmented volume color data of the VHP thorax were published by Zhou [18].

Other techniques for extracting visual information from volumes include isosurface extraction techniques such as the Marching Cubes algorithm [11]. This algorithm and its relatives have been used extensively to display the information within volume data, and have been successfully used on the Visible Human data. However, early work by Lorensen on the Visible Human data either used the X-ray CT data, or separated the red channel of the RGB images to generate isosurfaces from scalar rather than multi-valued data [10]. Similar work by the Vesalius group at Columbia University extracts isosurfaces from the color data, and later uses the original color volume as a solid texture to apply color information to the extracted polygonal surfaces [6]. Both of these approaches have involved strictly surface, rather than volume, rendering.

Researchers working with the Visible Human data have analyzed the color gamut of the male dataset in their work on photorealistic volume rendering and virtual dissections [8, 12]. Since reflectance of the light rays is often based on gradients measured in the object volume, careful consideration of the color spaces involved should be part of the work. Sapiro showed that selection of the color space can make dramatic differences when attempting nonlinear image processing [14].

4 Volume Shading

The determination of opacity values and other non-color material properties during the illumination process is a key challenge of direct volume rendering of photographic volumes. Although color information is available in the photographic description of each voxel, the data contains no information on the viewpoint dependent reflective or light transmission properties of the voxels.

The color values of each voxel describe the total reflection from that voxel for a specific set of lighting and viewing parameters. Specifically, the color shows the voxel lit and viewed from directly above. If we treat this total reflection as just diffuse reflection, we can use the same color as the reflection from other viewpoints, since diffuse reflection is independent of view direction. While using a correct bidirectional reflection distribution function for the voxel would yield more accurate results, determining the BRDF from photographic data is still an open research question. Considering the reflection to be simply diffuse results in a reasonable first approximation.

Determining the transmission properties, both opacity and attenuation properties, of each voxel is a larger problem. As with detailed reflection properties, this information is not available in the data. Unlike detailed reflection properties, it cannot simply be ignored. Rendering the volume with constant opacity yields unsatisfying results where little structure of the volume is clearly visible.

While using transmission properties that mimic those of the actual tissue types is attractive, segmentation from color data is particularly difficult. In particular, unlike the density values in CT data, there is no unique relation between color and tissue type.

Our initial approach has been to perform a variant of gradient-based shading on the photographic volume. Gradient-based shading emphasizes boundaries between regions with different properties, in this case different colors. Unlike threshold-based surface rendering approaches, however, gradient-based direct volume rendering still shows re-

gions of lesser gradient. Surface rendering shows only boundaries, while volume rendering can show the entire volume.

5 Implementation

Implementation considerations included the correct color space for gradient calculation and the specific gradient calculations themselves. We also found it useful to use some simple semi-automatic methods to extract segments of the data of particular interest. We chose to apply our methods on the brain section of the data since it represents a comparatively small anatomical volume and contains non-unique color attributes to challenge segmentation processing algorithms. A total of 139 sequential images contained portions of the brain.

5.1 Segment Extraction Preprocessing

In addition to visualizing the entire head, we wanted to experiment with visualizations of just the brain. To this end we used a sequence of semi-automatic filters to strip away unwanted tissue. First, voxels with colors unrepresentative of brain tissue were first removed. Essentially, this step left only voxels colored with a band of shades of yellow and brown. The HSV color model was used to specify desired colors, because of the lack of general intuition about the RGB components of flesh tones. This step stripped away much unwanted tissue, but left outliers in non-target regions that happened to match the specified color bands, as well as gaps within brain tissue in areas where the colors deviated from the specified bands. Next, a 3D flood fill process was used to select only those voxels meeting connectivity requirements with manually selected seed points. After this step, geographically isolated outliers have been eliminated, but internal gaps remain. Finally, an internal dilation filter is used to fill in interior gaps without dilating external image areas.

5.2 Color Space Transformation

RGB image slice values were converted to CIE l^*u^*v space to obtain a perceptually uniform representation of the color volume. A perceptually uniform color space has the characteristic that equal distances in the color space correspond to equal perceptual differences, at least for reasonably small distances. Using a perceptually uniform color space for gradient calculations allows us to emphasize those features which are noticeable in the photographs, creating a more realistic volume rendering than using a device-derived color space, such as RGB. The CIE l^*u^*v color space, in particular, also offers the advantage that chromatic and achromatic components of color are described by orthogonal color space dimensions. In the future, this feature will allow us to experiment with biased weighting of the chromatic and achromatic color components, just as the human visual system performs certain scene understanding tasks with segregated achromatic and chromatic color information.

Color space conversions were performed using the methods of Hall [4]. Because precise specifications of the color primaries of the image data were not directly available, we approximated them by the NTSC standard primaries. Although deviations of

actual primaries from these are expected to be modest, this approximation does not guarantee the correct absolute CIE l^*u^*v coordinates for voxels, compromising the device independence of the color space. However, the approximation should preserve the relationships between points in the color space, which is our primary concern since relative judgment, rather than absolute, are the basis for almost all perceptual processes. Specification of the image data color primaries is indirectly available from the color calibration card included in each photograph. Analysis of the RGB values for these physically measurable colors could provide the required calibration information. In fact, analysis of the calibration card colors could address another problem with the Visible Human data, that of inconstant color calibration from slice to slice. This problem presumably results primarily from changes in apparatus illumination over the course of the data collection process, and is apparent as dark and light stripes across coronal or sagittal slices. Since we know that the card did not undergo substantial changes in color over the course of data collection, color differences from card image to card image must represent changes in illumination and camera parameters. Ideally, a single process could be used to calibrate the image data to the color card, as well as to the CIE color primaries. We have not yet undertaken this process, but as better calibration data becomes available, the approximations made in the color model transformation could simply be corrected.

5.3 Rendering Process

We performed rendering experiments on both the unsegmented and the segmented brain photographic (rgb) data (139x600x500), which were converted to CIE l^*u^*v space. The volume gradient was approximated using a standard central difference approximation with a change to account for the distance of the colors in the CIE l^*u^*v space instead of using the voxel density differences:

$$\begin{aligned} \text{gradient}.x &= CIE_distance(\text{voxel}[x-1][y][z], \text{voxel}[x+1][y][z]) \\ \text{gradient}.y &= CIE_distance(\text{voxel}[x][y-1][z], \text{voxel}[x][y+1][z]) \\ \text{gradient}.z &= CIE_distance(\text{voxel}[x][y][z-1], \text{voxel}[x][y][z+1]) \end{aligned}$$

where

$$CIE_distance(\text{voxel1}, \text{voxel2}) = \sqrt{(\text{voxel1}.l - \text{voxel2}.l)^2 + (\text{voxel1}.u - \text{voxel2}.u)^2 + (\text{voxel1}.v - \text{voxel2}.v)^2}.$$

For the density of each voxel, the magnitude of the gradient was used. The volume renderer used is a modified volume ray tracer that uses atmospheric accumulation, attenuation, illumination, and shadowing [17, 3]. The opacity transfer function that was used is the following:

$$\text{opacity} = (\text{gradient_magnitude} * \text{scalar})^{\text{exponent}}$$

6 Results

The results demonstrating the effectiveness of using the CIE l^*u^*v color space distance as the basis for volume rendering from photographic data can be seen in Figures 2a and 2b. Figure 2a shows the brain rendered with a lower opacity scalar, 0.9, and an exponent of 1. Figure 2b shows a more opaque volume created by a higher opacity scalar, 1.0, and using an exponent less than 1 (specifically, 0.8) to increase the importance of small gradient changes within the volume. Both images highlight the tissue boundaries between the grey matter and the Corpus Collosum and also reveal a lateral ventricle within the brain.

Figure 3a, 3b, 3c, and 3d show volume rendering from the unsegmented photographic data. Figure 3a shows a side-view image generated by a higher exponent value (exponent=1.2) making the interior tissue structures more transparent. Figure 3b shows the same view image generated with a lower exponent (exponent=0.9) to increase the importance of gradients within tissue types. Both of these images show tissue boundaries between grey matter and the Corpus Collosum, a lateral ventricle, portions of the skull, and the sinus cavity. Figures 3c and 3d are top view images of the unsegmented photographic data showing the range of images that can be generated by varying the exponent to show all internal tissues (Figure 3c, exponent=0.5) or to show primarily tissue boundaries (Figure 3d, exponent=1.0).

7 Summary and Future Directions

We have proposed a basic framework for addressing the rendering challenges presented by photographic volume data. Preliminary results show that color gradient in a perceptually uniform color space can be used to perform gradient-based volume rendering from photographic data.

Many open issues remain. These include color calibration and internal registration methods, the optimal weighting of color space dimensions in gradient calculations, more realistic reflectance models, and determination of segment-based densities directly from the photographic volume.

8 Acknowledgments

This work supported in part by the National Science Foundation under Grants No. ACIR 9996043 and ACIR 9978032.

References

1. B. Cabral, N. Cam, and J. Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *1994 Symposium on Volume Visualization*, pages 91–98, 1994.
2. R.A. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. In *Computer Graphics (SIGGRAPH '88 Proceedings)*, number 22, pages 65–74, 1988.

3. David S. Ebert and Richard E. Parent. Rendering and animation of gaseous phenomena by combining fast volume and scanline A-buffer techniques. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 357–366, August 1990.
4. Roy Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, 1989.
5. T. Hazeldine. Annular slicing: Physics, materials, geology, sedimentology and medicine make use of precision annular cutting, so does the semiconductor industry. *Semiconductor*, Oct. 1997.
6. C. Imielinska. Technical challenges of 3d visualization of large color data sets. In R. A. Banvard and P. Cerveri, editors, *Proceedings of the Second Visible Human Project Conference*. US Department of Health and Human Services, Public Health Service, National Institutes of Health, October 7-8 1998.
7. Arie Kaufman. Volume visualization. *ACM Computing Surveys*, 28(1):165–167, 1996.
8. J. Kerr. Photorealistic volume-rendered anatomical atlases and interactive virtual dissections of the dissectible human. In R. A. Banvard, editor, *Proceedings of the Visible Human Project Conference*. US Department of Health and Human Services, Public Health Service, National Institutes of Health, October 7-8 1996.
9. Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, 1988.
10. William E. Lorensen. Marching through the visible man. In *Proceedings of IEEE Visualization '95*, pages 368–373. IEEE Press, October 1995.
11. William E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 163–169, 1987.
12. J. Marquez. Radiometric inhomogeneities in the color cryosection images of the vhp. In R. A. Banvard, editor, *Proceedings of the Visible Human Project Conference*. US Department of Health and Human Services, Public Health Service, National Institutes of Health, October 1-2 1996.
13. W. Nip and C. Logan. Whole frog technical report. Technical Report LBL-35331, University of California, Lawrence Berkeley Laboratory, 1991.
14. G. Sapiro and D. Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Transactions of Image Processing*, 5:1582–1586, 1996.
15. S. Spitzer, M. J. Ackerman, A. L. Scherzinger, and D. Whitlock. The visible human male: A technical report. *Journal of the American Medical Informatics Association*, 3(2):118–130, 1996.
16. A. W. Toga and J. C. Mazziotta. *Brain Mapping: The Methods*. Academic Press, 1996.
17. Roni Yagel, David S. Ebert, J. N. Scott, and Y. Kurzion. Grouping volume renderers for enhanced visualization in computational fluid dynamics. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):117–132, June 1995. ISSN 1077-2626.
18. Ruixa Zhou and Earl Henderson. Visualization of the visible human anatomical images. In R. A. Banvard and P. Cerveri, editors, *Proceedings of the Second Visible Human Project Conference*. US Department of Health and Human Services, Public Health Service, National Institutes of Health, October 1-2 1998.



Fig. 1. Sample slice of the digital photographic data transecting the eyes and optic nerves of the Visible Human Male Dataset.




Fig. 2. Volume rendering of segmented photographic Visible Male brain, without and with enhancement of low gradient areas.




Fig. 3. Volume rendering of photographic Visible Male head, without and with enhancement of low gradient areas from side and top views.