Expressive Volume Rendering

Penny Rheingans
University of Maryland Baltimore

1000 Hilltop Circle

21250, Baltimore, MD USA
rheingans@cs.umbc.edu

ABSTRACT

Accurately and automatically conveying the structure of a volume model is a problem not fully solved by existing volume rendering approaches. Physics-based volume rendering approaches create images that may match the appearance of translucent materials in nature, but may not embody important structural details. Transfer function approaches allow flexible design of the volume appearance, but generally require substantial hand tuning for each new data set in order to be effective. We have introduced the volume illustration approach, combining the familiarity of a physics-based illumination model with the ability to enhance important features using non-photorealistic rendering techniques. Since features to be enhanced are defined on the basis of higher-order model characteristics rather than volume sample value, the application of volume illustration techniques requires less manual tuning than the design of a good transfer function. Volume illustration provides a flexible unified framework for enhancing structural perception of volume models through the amplification of features, the addition of illumination effects, and the application of procedural textures. Volume illustration works on both presampled and procedurally defined volume models, enabling a range of image styles from practical technical illustrations to more abstract painterly effects.

Keywords

Direct volume rendering, scientific illustration, non-photorealistic rendering, data visualization.

1. INTRODUCTION

The main goal of visualization is to effectively convey information to the user using the wide input channel provided by the human visual system. A visual representation of a large data set can capture both interesting elements and interesting structure in the data. For volume data, one key goal is to convey the structure of the data distribution, for example the shape of the liver, or the extent of ground water contamination, or the density of ozone throughout the atmosphere. In some applications, the boundaries of interesting regions are sharp and well-defined, for instance the boundaries of organs or the places where groundwater contamination exceeds legal limits. In other applications, the boundaries of interesting regions are rather diffuse, for instance the boundaries of tumors or molecules. And in still other applications, the whole structure of the data distribution is of interest, rather than just the boundaries of particular regions. Effective general volume visualization techniques must address all three kinds of applications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Journal of WSCG, Vol.12, No.1-3, ISSN 1213-6972
WSCG'2004, February 2-6, 2004, Plzen, Czech Republic.
Copyright UNION Agency – Science Press

Algorithms for the visualization of volume data can be categorized into two general approaches. Surface algorithms first map the volume data to representative geometry, such as an isosurface of constant value, and then render the geometric representation using standard rendering techniques. The second type of approach, direct volume rendering, generates the image directly from the volume data, without first creating any geometry. For volume models, the key advantage of direct volume rendering over surface rendering approaches is the potential to show the structure of the value distribution throughout the volume, rather than just at selected boundary surfaces of variable value (by isosurface) or coordinate value (by cutting plane). The contribution of each volume sample to the final image is explicitly computed and included. The key challenge of direct volume rendering is to convey that value distribution clearly and accurately. In particular, showing each volume sample with sufficient clarity and opacity that its structure is apparent but not so much that volume samples in the rear of the volume are overly obscured.

Volume illustration is a new approach to volume rendering involving the augmentation of a physics-based rendering process with non-photorealistic rendering (NPR) techniques to enhance the expressiveness of the visualization [Ebe00, Rhe01,Lu03]. NPR draws inspiration from such fields as art and technical illustration to develop automatic methods to synthesize images with an illustrated look from geometric surface models. Non-photorealistic rendering research has effectively addressed both the illustration of surface shape and the visualization of 2D data, but has virtually ignored the rendering of volume models. Volume illustration introduces a set of NPR techniques specifically for the visualization of volume data, including both the adaptation of existing NPR techniques to volume rendering and the development of new techniques specifically suited for volume models.

2. MOTIVATION

Volumetric illustration differs from surface-based NPR in several important ways. In surface-based NPR, the surfaces (features) are well defined, whereas with volumes, the volumetric features vary continuously throughout three-dimensional space and are not as well defined as surface features. Feature areas within the volume must be determined through analysis of local volumetric properties. Once these volumetric feature volumes are identified, user selected parametric properties can be used to enhance and illustrate them.

In a surface model, the essential feature is the surface itself. The surface is explicitly and discretely defined by a surface model, making "surfaceness" a boolean quality. Many other features, such as silhouettes or regions of high curvature, are simply interesting parts of the surface. Such features can be identified by analysis of regions of the surface. In a volume model, there are no such discretely defined features. Additional processing is required to first identify interesting features in the volume.

Another difficulty with volumetric models is that few of the usual depth cues are present in traditional rendering of translucent volumes. Obscuration cues are largely missing since there are no opaque objects to show a clear depth ordering. Perspective cues from converging lines and texture compression are also lacking, since few volume models contain straight lines or uniform textures. The dearth of clear depth cues makes understanding spatial relationships of features in the volume difficult. One common approach to this difficulty is the use of hard transfer functions, those with rapidly increasing opacity at particular value ranges of interest. While this may increase depth cues by creating the appearance of surfaces within the volume, it does so by hiding all information in some regions of the volume, sacrificing a key advantage of volume rendering.

Similarly, information about the orientation of features within the volume is also largely missing. Transfer function approaches typically perform no illumination calculations to determine the color at a volume sample. Nor does the transfer function approach usually include the effect of shadows, particularly self-shadows in which the volume shadows itself. Although physics-based volume rendering algorithms do include illumination and shadow effects, the effects are generally subtle and difficult to interpret unambiguously. As a result, the shape of individual structures within the volume is difficult to perceive, as can be seen in the volume rendering of an abdominal CT volume in Figure 1a.

3. KEY ILLUSTRATION GOALS

Effective illustrations are not merely faithful reproductions of the physical form of the subject. A skilled illustrator creates a more effective result by deviating from the exact appearance in specific intentional ways. It is these deviations from exact reality that can make the illustration more expressive than a photograph would be. More information about scientific illustration can be found in texts on illustration and various application areas (including Woo94,Cla99, Hod03).

Four specific goals of illustration are particularly interesting. These goals are the desire to direct attention to particular parts of the subject, to clearly convey the shape of the subject, to emphasize depth relationships of items in the scene, and to convey the nature of translucence of objects. This last goal, while not particularly common in many types of illustration, is of particular relevance to those interested in capturing volumetric entities.

One key technique illustrators use to direct attention is selective simplification. Areas of lower interest are simplified to show gross structure for context, while areas that should receive greater attention are rendered in greater detail. Another technique for directing attention is that of vignetting, where regions outside the center of interest are made to appear to fade away to nothing.

Representation of the shape of the subject begins with the capture of the overall shape of the object. Illustrators advise that squinting at an object removes detail to reveal the main contour shape. Additional aids to the representation of shape include believable lighting and shadows. One particular form of stylized lighting effect is that of a rim shadow. A *rim shadow* shows the periphery and distant areas of the subject to be in shadow, simulating a beam of light from the front. *Tone shading* adds warm/cool cues to shape, simulating the chromatic warming produced by a warm-specturm light source. Outlines serve an important purpose in conveying shape by emphasizing important boundaries and providing detail in flat parts that would not be captured by lighting cues.

Looking into the far distance, particularly in hazy conditions, shows distant features as having less contrast and detail at they fade into the haze. Painters have used forms of this *aerial perspective* for at least as far back as Leonardo da Vinci [daV1506]. Painterly aerial perspective may also include blue tones in distant objects, capitalizing (perhaps unintentionally) on the wave-length dependent refraction effects in the lens of the human eye. Illustrators exaggerate this atmospheric perspective effect to show depth relationships over much smaller distances. Another technique for clarifying depth relationships in an illustration is the inclusion of halos around foreground objects to make them stand out from objects behind.

Objects which are themselves translucent are generally represented in illustration through changes in the appearance of objects seen through them. Specifically, the appearance of underlying objects is muted and contains less detail. A related technique shows the location of objects inside other objects that are not naturally translucent. In this technique, the appearance of the entering object is strengthened before it enters another, eliminated entirely upon entering, and then faded back in.

4. NPR VISUALIZATION

There has been extensive research for illustrating surface shape using non-photorealistic rendering techniques [for an overview, see Goo01, Str02]. A few researchers have applied NPR techniques to the display of data. Laidlaw used concepts from painting to create visualizations of 2D data, using brushstroke-like elements to convey information [Lai98] and a painterly process to compose complex visualizations [Kir99]. Treavett has developed techniques for pen-and-ink illustrations of surfaces within volumes [Tre00].

The key difficulty of adapting surface techniques for NPR is that there are no explicit surfaces in volume models. So, instead of simply calculating illustration appearance on surfaces, volume illustration must first explicitly identify features to be illustrated and then render these features appropriately. These features may be either boolean (as surfaces are) or may be continuously defined to indicate the strength of a feature at a location. Continuously defined features, in particular, generally require the consideration of each voxel in the rendering process.

Although volume data contains no surfaces, the boundaries between regions may still be of interest. The local gradient magnitude at a volume sample can be used to indicate the degree to which the volume sample is a boundary between disparate regions. The direction of the gradient is analogous to the surface normal. Regions of high gradient are similar to surfaces, but now “surfaceness” is a continuous, rather than boolean, quality. This makes gradient a straight-forward feature indicator.

Some NPR techniques originally developed for surfaces can be adapted for volumes including a boolean or continuous feature indicator in the appearance model, resulting in strong illustration where there are strong features and fainter illustration for fainter features. In this way, the feature strength scales the illustration weight by modifying voxel opacity, stroke density, or other rendering parameters. The techniques most suitable for this adaption are those which can be characterized by some appearance model evaluated at sample points (for instance, the tone shading model of Gooch et al. [Goo98] or the 3D textures of Treavett and Chen [Tre00]) or those in which strokes of differing densities are generated in image space (for instance, Deu00).

More recent work has specifically addressed the non-photorealistic rendering of unthresholded volume data, both in terms of new techniques [Ebe00, Tre01, Rhe01, Lu03] and performance enhancements [Cse01, Lu02, Lum02].

5. ACKNOWLEDGMENTS

From the beginning, this work has been a collaboration with David Ebert of Purdue University. A number of others have contributed to the ideas expressed here, especially Christopher Morris and Aidong Lu. This material is based upon work

Figure 1. Three renderings of an abdominal CT data set. a) standard volume rendering, b) volume illustration with boundary and silhouette enhancement, c) volume illustration with tone shading.

supported by the National Science Foundation under Grant No. ACIR 0081581.

6. REFERENCES

[Cla99] John O,. E. Clark, *A Visual Guide to the Human Body*, Barnes and Noble Books, 1999.

[Cse01] Balazs Csébfalvi, Lukas Mroz, Helwig Hauser, Andreas König, Eduard Gröller. Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum*, vol. 20, no. 3, pp. 452-460, 2001.

[daV1506] Leonardo daVinci, *The Virgin of the Rocks*, National Gallery, London, 1503-1506.

[Due00] O. Duessen, S. Hiller, C. van Overveld, and T. Strothotte. "Floating Points: A Method for Computing Stipple Drawings". *Proceedings of Eurographics 2000*. vol. 19. no. 3. The Eurographics Association and Blackwell Publishers. Malden, MA. 2000.

[Ebe00] David Ebert and Penny Rheingans, Volume Illustration: Non-Photorealistic Rendering of Volume Models, Proceedings of IEEE Visualization 2000, pp. 195-202, October 2000.

[Goo98] Amy Gooch, Bruce Gooch, Peter Shirley, and Elaine Cohen. A Non-photorealistic Lighting Model for Automatic Technical Illustration. In *Proceedings of SIGGRAPH '98* (Orlando, FL, July 1998), Computer Graphics Proceedings, Annual Conference Series, pp. 447-452, July 1998.

[Goo01] Bruce and Amy Gooch. *Non-Photorealistic Rendering*. A.K. Peters, 2001.

[Gri94] Griffel, Lois, Painting the Impressionist Landscape, Watson-Guptill Publications,1994.

[Hod03] Elaine Hodges, ed.. *The Guild Handbood of Scientific Illustration*, second edition, John Wiley and Sons, Inc., 2003.

[Kir99] R.M. Kirby, H. Marmanis, and D.H. Laidlaw. Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting, *IEEE Visualization '99*, pp. 333-340, 1999.

[Lai98] David H. Laidlaw, Eric T. Ahrens, David Kremers, Matthew J. Avalos, Russell E. Jacobs, and Carol Readhead. Visualizing Diffusion Tensor Images of the Mouse Spinal Cord, *IEEE Visualization '98*, pp. 127-134, 1998.

[Lu02] Aidong Lu, Joe Taylor, Mark Harner, David Ebert and Charles Hansen. Hardware-Accelerated Interactive Illustrative Stipple Drawing of Polygonal Objects, *Proceedings of VMV2002: Vision, Modeling, and Visualization*, pp. 61-68, 2002.

[Lu03] Aidong Lu, Christopher Morris, Joe Taylor, David Ebert, Charles Hansen, Penny Rheingans, and Mark Hartner. Illustrative Interactive Stipple Rendering, *IEEE Transactions on Visualization and Computer Graphics*, vol.9, no. 2, pp. 127-138, April-June 2003.

[Lum02] Eric Lum and Kwan-liu Ma, Hardware-Accelerated Parallel Non-Photorealistic Volume Rendering, *Proceedings of the Second International Symposium on Non-Photorealistic Animation and Rendering*, pp. 67-, 2002.

[Rhe01] Penny Rheingans and David Ebert. Volume Illustration: Nonphotorealistic Rendering of Volume Models, *IEEE Transactions on Visualization and Computer Graphics*, vol. 7, no. 3, pp.253-264, July-September 2001.

[Str02] Thomas Strothotte and Stefan Schlechtweg. *Non-Photorealistic Computer Graphics: Modeling, Rendering, and Animation*, Morgan Kaufmann, 2002.

[Tre00] S.M.F. Treavett and M. Chen. Pen-and-Ink Rendering in Volume Visualisation, *Proceedings of IEEE Visualization 2000,* pp. 203-209, October 2000.

[Tre01] S.M.F. Treavett, M. Chen, R. Satherley, and M.W. Jones, Volumes of Expression: artistic modelling and rendering of volume data sets, to appear in *Proceedings of Computer Graphics International (CGI2001)*, Hong Kong, July 2001.

[Woo94] Phyllis Wood. *Scientific Illustration*, second edition, John Wiley and Sons, Inc., 1994.