

Volume Illustration: Non-Photorealistic Rendering of Volume Models

David Ebert Penny Rheingans
Computer Science and Electrical Engineering
University of Maryland Baltimore County
Baltimore MD 21250
[ebert | rheingan]@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 which 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 introduce 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 local volume 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 and the addition of illumination effects.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism – color, shading, and texture; I.3.8 [Computer Graphics]: Applications.

Keywords: Volume rendering, non-photorealistic rendering, illustration, lighting models, shading, visualization.

1 Introduction

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 full opacity and clarity is impossible if volume samples in the rear of the volume are not to be completely obscured.

Traditionally, volume rendering has employed one of two approaches. The first attempts a physically accurate simulation of a process such as the illumination and attenuation of light in a gaseous volume or the attenuation of X-rays through tissue [Kajiya84, Drebin88]. This approach produces the most realistic

and familiar views of a volume data set, at least for data that has an appropriate physical meaning. The second approach is only loosely based on the physical behavior of light through a volume, using instead an arbitrary transfer function specifying the appearance of a volume sample based on its value and an accumulation process that is not necessarily based on any actual accumulation mechanism [Levoy90]. This approach allows the designer to create a wider range of appearances for the volume in the visualization, but sacrifices the familiarity and ease of interpretation of the more physics-based approach.

We propose a new approach to volume rendering: the augmentation of a physics-based rendering process with non-photorealistic rendering (NPR) techniques [Winkenbach94, Salisbury94] to enhance the expressiveness of the visualization. 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. We describe 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. We call this approach *volume illustration*.

The volume illustration approach combines the benefits of the two traditional volume rendering approaches in a flexible and parameterized manner. It provides the ease of interpretation resulting from familiar physics-based illumination and accumulation processes with the flexibility of the transfer function approach. In addition, volume illustration provides flexibility beyond that of the traditional transfer function, including the capabilities of local and global distribution analysis, and light and view direction specific effects. Therefore, volume illustration techniques can be used to create visualizations of volume data that are more effective at conveying the structure within the volume than either of the traditional approaches. As the name suggests, volume illustration is intended primarily for illustration or presentation situations, such as figures in textbooks, scientific articles, and educational video.

2 Related Work

Traditional volume rendering spans a spectrum from the accurate to the ad hoc. Kajiya's original work on volume ray tracing for generating images of clouds [Kajiya84] incorporated a physics-based illumination and atmospheric attenuation model. This work in realistic volume rendering techniques has been extended by numerous researchers [Nishita87, Ebert90, Krueger91, Williams92, Max95, Nishita98]. In contrast, traditional volume rendering has relied on the use of transfer functions to produce artificial views of the data to highlight regions of interest [Drebin88]. These transfer functions, however, require in-depth knowledge of the data and need to be adjusted for each data set.

The design of effective transfer functions is still an active research area [Fang98, Kindlmann98, Fujishiro99]. While transfer functions can be effective at bringing out the structure in the value distribution of a volume, they are limited by their dependence on voxel value as the sole transfer function domain.

In contrast, there has been extensive research for illustrating surface shape using non-photorealistic rendering techniques. Adopting a technique found in painting, Gooch et al. developed a tone-based illumination model that determined hue, as well as intensity, from the orientation of a surface element to a light source [Gooch98]. The extraction and rendering of silhouettes and other expressive lines has been addressed by several researchers [Saito90, Salisbury94, Gooch99, Interrante95]. Expressive textures have been applied to surfaces to convey surface shape [Rheingans96, Salisbury97, Interrante97].

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 [Laidlaw98] and a painterly process to compose complex visualizations [Kirby99]. Treavett has developed techniques for pen-and-ink illustrations of surfaces within volumes [Treavett00]. Interrante applied principles from technical illustration to convey depth relationships with halos around foreground features in flow data [Interrante98]. Saito converted 3D scalar fields into a sampled point representation and visualized selected points with a simple primitive, creating an NPR look [Saito94]. With the exceptions of the work of Saito and Interrante, the use of NPR techniques has been confined to surface rendering.

3 Approach

We have developed a collection of volume illustration techniques that adapt and extend NPR techniques to volume objects. Most traditional volume enhancement has relied on functions of the volume sample values (e.g., opacity transfer functions), although some techniques have also used the volume gradient (e.g., [Levoy90]). In contrast, our volume illustration techniques are fully incorporated into the volume rendering process, utilizing viewing information, lighting information, and additional volumetric properties to provide a powerful, easily extensible framework for volumetric enhancement. Comparing Diagram 1, the traditional volume rendering system, and Diagram 2, our volume illustration rendering system, demonstrates the difference in our approach to volume enhancement. By incorporating the enhancement of the volume sample's color, illumination, and opacity into the rendering system, we can implement a wide range of enhancement techniques. The properties that can be incorporated into the volume illustration procedures include the following:

- Volume sample location and value
- Local volumetric properties, such as gradient and minimal change direction
- View direction
- Light information

The view direction and light information allows global orientation information to be used in enhancing local volumetric features. Combining this rendering information with user selected parameters provides a powerful framework for volumetric enhancement and modification for artistic effects.

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

We begin with a volume renderer that implements physics-based illumination of gaseous phenomena. The opacity transfer function that we are using is the following simple power function:

$$o_v = (k_{os} v_i)^{k_{oe}}$$

where v_i is the volume sample value and k_{os} is the scalar controlling maximum opacity. Exponent k_{oe} values less than 1 soften volume differences and values greater than 1 increase the contrast within the volume.

Figure 1 shows gaseous illumination of an abdominal CT volume of $256 \times 256 \times 128$ voxels. In this image, as in others of this dataset, the scene is illuminated by a single light above the volume and slightly toward the viewer. The structure of tissues and organs is difficult to understand. In Figure 2, a transfer function has been used to assign voxel colors which mimic those found in actual tissue. The volume is illuminated as before. Organization of tissues into organs is clear, but the interiors of structures are still unclear. We chose to base our examples on an atmospheric illumination model, but the same approach can be easily applied to a base renderer using Phong illumination and linear accumulation.

In the following two sections, we describe our current collection of volume illustration techniques. These techniques can be applied in almost arbitrary amounts and combinations, becoming a flexible toolkit for the production of expressive images of volume models. The volume illustration techniques we

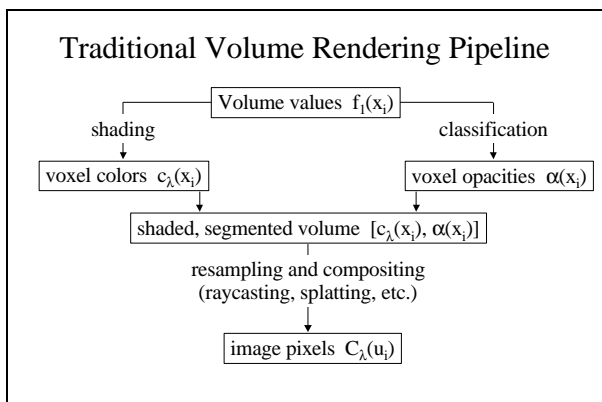


Diagram 1. Traditional Volume Rendering Pipeline.

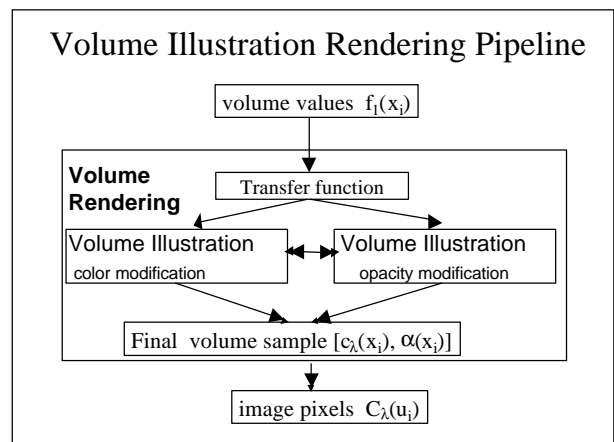


Diagram 2. Volume Illustration Rendering Pipeline.

To appear in *Proceedings of IEEE Visualization '00* (October 2000, Salt Lake City, UT).

propose are of two basic types: feature enhancement, and depth and orientation cues.

4 Feature Enhancement

In a surface model, the essential feature is the surface itself. The surface is explicitly and discretely defined by a surface model, making “surfacedness” 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. Volume characteristics and the features that they indicate exist continuously throughout the volume. However, the boundaries between regions are still one feature of interest. The local gradient magnitude at a volume sample can be used to indicate the degree to which the 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 “surfacedness” is a continuous, volumetric quality, rather than a boolean quality. We have developed several volume illustration techniques for the enhancement of volume features based on volume gradient information.

4.1 Boundary Enhancement

Levoy [Levoy90] introduced gradient-based shading and opacity enhancement to volume rendering. In his approach, the opacity of each voxel was scaled by the voxel's gradient magnitude to emphasize the boundaries between data (e.g., tissue) of different densities and make areas of constant density transparent (e.g., organ interiors). We have adapted this idea to allow the user to selectively enhance the density of each volume sample by a function of the gradient. Assume a volume data set containing a precomputed set of sample points. The value at a location P_i is a scalar given by:

$$v_i = f(P_i) = f(x_i, y_i, z_i)$$

We can also calculate the value gradient $\tilde{N}_f(P_i)$ at that location. In many operations we will want that gradient to be normalized. We use \tilde{N}_{fn} to indicate the normalized value gradient.

Before enhancement, voxel values are optionally mapped through a standard transfer function which yields color value c_v and opacity o_v for the voxel. If no transfer function is used, these values can be set to constants for the whole volume. The inclusion of a transfer function allows artistic enhancements to supplement, rather than replace, existing volume rendering mechanisms.

We can define a boundary-enhanced opacity for a voxel by combining a fraction of the voxel's original opacity with an enhancement based on the local boundary strength, as indicated by the voxel gradient magnitude. The gradient-based opacity of the volume sample becomes:

$$o_g = o_v \left(k_{gc} + k_{gs} \left(\|\nabla_f\| \right)^{k_{ge}} \right)$$

where o_v is original opacity and \tilde{N}_f is the value gradient of the volume at the sample under consideration. This equation allows the user to select a range of effects from no gradient enhancement ($k_{gc}=1, k_{gs}=0$) to full gradient enhancement ($k_{gs} >=1$) to only showing areas with large gradients ($k_{gc}=0$), as in traditional volume rendering. The use of the power function with exponent k_{ge} allows the user to adjust the slope of the opacity curve to best highlight the dataset.

Figure 3 shows the effect of boundary enhancement in the medical volume. The edges of the lungs and pulmonary vasculature can be seen much more clearly than before, as well as some of the internal structure of the kidney. Parameter values used in Figure 3 are $k_{gc} = 0.7, k_{gs} = 10, k_{ge} = 2.0$.

4.2 Oriented Feature Enhancement: Silhouettes, Fading, and Sketch Lines

Surface orientation is an important visual cue that has been successfully conveyed by artists for centuries through numerous techniques, including silhouette lines and orientation-determined saturation effects. Silhouette lines are particularly important in the perception of surface shape, and have been utilized in surface illustration and surface visualization rendering [Salisbury94, Interrante95]. Similarly, silhouette volumes increase the perception of volumetric features.

In order to strengthen the cues provided by silhouette volumes, we increase the opacity of volume samples where the gradient nears perpendicular to the view direction, indicated by a dot product between gradient and view direction which nears zero. The silhouette enhancement is described by:

$$o_s = o_v \left(k_{sc} + k_{ss} \left(1 - \text{abs}(\nabla_{fn} \cdot V) \right)^{k_{se}} \right)$$

where k_{sc} controls the scaling of non-silhouette regions, k_{ss} controls the amount of silhouette enhancement, and k_{se} controls the sharpness of the silhouette curve.

Figure 4 shows the result of both boundary and silhouette enhancement in the medical volume. The fine honeycomb structure of the liver interior is clearly apparent, as well as additional internal structure of the kidneys. Parameter values used in Figure 4 are $k_{gc} = 0.8, k_{gs} = 5.0, k_{ge} = 1.0; k_{sc} = 0.9, k_{ss} = 50, k_{se} = 0.25$.

Decreasing the opacity of volume features oriented toward the viewer emphasizes feature orientation, and in the extreme cases, can create sketches of the volume, as illustrated in Figure 5. Figure 5 shows a black and white sketch of the medical dataset by using a white sketch color and making non volumetric silhouettes transparent. To get appropriate shadowing of the sketch lines, the shadows are calculated based on the original volume opacity. Using a black silhouette color can also be effective for outlining volume data.

Orientation information can also be used effectively to change feature color. For instance, in medical illustration the portions of anatomical structures oriented toward the viewer are desaturated and structures oriented away from the view are darkened and saturated [Clark99]. We simulate these effects by allowing the volumetric gradient orientation to the viewer to modify the color, saturation, value, and transparency of the given volume sample. The use of the HSV color space allows the system to easily utilize the intuitive color modification techniques of painters and illustrators. Figure 10 shows oriented changes in the saturation and value of the medical volume. In this figure, the color value (V) is decreased as the angle between the gradient and the viewer increases, simulating more traditional illustration techniques of oriented fading.

5 Depth and Orientation Cues

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. Many volume rendering systems use very simple illumination models and often do not include the effect of shadows, particularly volume self-shadowing to improve performance, even though many volume shadowing algorithms have been developed [Ebert90, Kajiyama84]. Accurate volumetric shadowing often produces subtle effects which do not provide strong three-dimensional depth cues. As a result, the shape of individual structures within even illuminated volumes is difficult to perceive.

We have developed several techniques for the enhancement of depth and orientation cues in volume models, inspired by shading concepts in art and technical illustration.

5.1 Distance color blending

Intensity depth-cuing is a well known technique for enhancing the perception of depth in a scene [Foley96]. This technique dims the color of objects far from the viewer, creating an effect similar to viewing the scene through haze. We have adapted this technique for volume rendering, dimming volume sample colors as they recede from the viewer. In addition, we have augmented the standard intensity depth-cuing with a subtle color modulation. This color modulation increases the amount of blue in the colors of more distant volume samples, simulating techniques used for centuries by painters, such as aerial perspective [daVinci1506, Beirstadt1881]. This technique exploits the tendency of cool colors (such as blue) to recede visually while warm colors (such as red) advance.

Depth-cued colors start as the voxel color at the front of the volume, decreasing in intensity and moving toward the background color as depth into the volume increases. The progression of depth-cuing need not be linear; we use an exponential function to control the application of depth-cuing. The distance color blending process can be described by:

$$c_d = (1 - k_{ds} d_v^{k_{de}}) c_v + k_{ds} d_v^{k_{de}} c_b$$

where k_{ds} controls the size of the color blending effect, k_{de} controls the rate of application of color blending, d_v is the fraction of distance through the volume, and c_b is a defined background color. When c_b is a shade of grey ($c_b = (a, a, a)$ for some value of a), only standard intensity depth-cuing is performed. Using a background color that is a shade of blue ($c_b = (a, b, c)$ for $c > a, b$), introduces a cool shift in distant regions. Other color modulation effects are clearly possible, but make less sense perceptually.

Figure 6 shows the effect of distance color blending. The ribs behind the lungs fade into the distance and the region around the kidneys seems to recede slightly. Color blending parameters used in Figure 6 are $c_b = (0, 0, 0.15)$, $k_{ds} = 1.0$, $k_{se} = 0.5$.

5.2 Feature halos

Illustrators sometimes use null halos around foreground features to reinforce the perception of depth relationships within a scene. The effect is to leave the areas just outside surfaces empty, even if an accurate depiction would show a background object in that place. Interrante [Interrante98] used a similar idea to show depth

relationships in 3D flow data using Line Integral Convolution (LIC). She created a second LIC volume with a larger element size, using this second volume to impede the view. Special care was required to keep objects from being obscured by their own halos. The resulting halos achieved the desired effect, but the method depended on having flow data suitable for processing with LIC.

We introduce a more general method for creating halo effects during the illumination process using the local spatial properties of the volume. Halos are created primarily in planes orthogonal to the view vector by making regions just outside features darker and more opaque, obscuring background elements which would otherwise be visible. The strongest halos are created in empty regions just outside (in the plane perpendicular to the view direction) of a strong feature.

The halo effect at a voxel is computed from the distance weighted sum of haloing influences in a specified neighborhood. In order to restrict halos to less interesting regions, summed influences are weighted by the complement of the voxel's gradient. The size of the halo effect is given by:

$$h_i = \left(\sum_n^{\text{neighbors}} \frac{h_n}{\|P_i - P_n\|^2} \right) (1 - \|\nabla_f(P_i)\|)$$

where h_n is the maximum potential halo contribution of a neighbor. The haloing influence of a neighbor is inversely related to its distance and the tendency of a location to be a halo is inversely related to its gradient magnitude.

The maximum potential halo contribution of each neighbor is proportional to the product of the alignment of the neighbor's gradient with the direction to the voxel under consideration (calculated from the dot product between them) and the degree to which the neighbor's gradient is aligned perpendicular to the view direction (also calculated as a dot product). The halo potential (h_n) is given by:

$$h_n = \left(\nabla_{fn}(P_n) \cdot \left(\frac{P_i - P_n}{\|P_i - P_n\|} \right) \right)^{k_{hpe}} (1 - \nabla_{fn}(P_n) \cdot V)^{k_{hse}}$$

where k_{hpe} controls how directly the neighbor's gradient must be oriented toward the current location, and k_{hse} controls how tightly halos are kept in the plane orthogonal to the view direction. The most strong halo effects will come from neighbors that are displaced from the volume sample of interest in a direction orthogonal to the view direction and that have a large gradient in the direction of this sample.

Once the size of the halo effect has been determined, parameters control the visual appearance of halo regions. The most common adjustment to the halo region is to decrease the brightness by a scalar times the halo effect and increase the opacity by another scalar times the halo effect. This method produces effects similar to those of Interrante, but can be applied to any type of data or model during the illumination process. Since the halos generated are inherently view dependent, no special processing must be done to keep features from casting a halo on themselves.

Figure 6 shows the effectiveness of adding halos to the medical volume. Structures in the foreground, such as the liver and kidneys, stand out more clearly. Halo parameters used in Figure 6 are $k_{hpe} = 1.0$ and $k_{hse} = 2.0$.

5.3 Tone shading

Another illustrative technique used by painters is to modify the tone of an object based on the orientation of that object relative to

the light. This technique can be used to give surfaces facing the light a warm cast while surfaces not facing the light get a cool cast, giving effects suggestive of illumination by a warm light source, such as sunlight. Gooch et al. proposed an illumination model based on this technique [Gooch98], defining a parameterized model for effects from pure tone shading to pure illuminated object color. The parameters define a warm color by combining yellow and the scaled fully illuminated object color. Similarly, a cool color combines blue and the scaled ambient illuminated object color. The final surface color is formed by interpolation between the warm and cool color based on the signed dot product between the surface normal and light vector. The model assumes a single light source, generally located above the scene.

We implemented an illumination model similar to Gooch tone shading for use with volume models. As with Gooch tone shading, the tone contribution is formed by interpolation between the warm and cool colors based on the signed dot product between the volume sample gradient and the light vector. Unlike Gooch tone shading, the illuminated object contribution is calculated using only the positive dot product, becoming zero at orientations orthogonal to the light vector. This more closely matches familiar diffuse illumination models.

The color at a voxel is a weighted sum of the illuminated gaseous color (including any traditional transfer function calculations) and the total tone and directed shading from all directed light sources. The new tone illumination model is given by:

$$c = k_{ta} I_G + \sum_i^{N_L} (I_t + k_{td} I_o)$$

where k_{ta} controls the amount of gaseous illumination (I_G) included, N_L is the number of lights, k_{td} controls the amount of directed illumination included, I_t is the tone contribution to volume sample color, and I_o is the illuminated object color contribution. Although this model allows for multiple light sources, more than a few is likely to result in confusing images, since we are not used to interpreting complex illumination coming from many lights.

The tone contribution from a single light source is interpolated from the warm and cool colors, depending on the angle between the light vector and the sample gradient. It is given by:

$$I_t = \left((1.0 + \nabla_{fn} \cdot L) / 2 \right) c_w + \left(1 - (1.0 + \nabla_{fn} \cdot L) / 2 \right) c_c$$

where L is the unit vector in the direction of the light and

$$c_w = (k_{ty}, k_{ty}, 0), \quad c_c = (0, 0, k_{tb})$$

describe the warm and cool tone colors. Samples oriented toward the light become more like the warm color while samples oriented away from the light become more like the cool color.

The directed illumination component is related to the angle between the voxel gradient and the light direction, for angles up to 90 degrees. It is given by:

$$I_o = \begin{cases} k_{td} I_i (\nabla_{fn} \cdot L) : \nabla_{fn} \cdot L > 0 \\ 0 : \nabla_{fn} \cdot L \leq 0 \end{cases}$$

where k_{td} controls how much directed illumination is added.

Figure 7 shows modified tone shading applied to the uncolored medical volume. The small structure of the liver shows clearly, as does the larger structures of the kidney. The bulges of intestine at the lower right are much more clearly rounded 3D shapes than with just boundary and silhouette enhancement (Figure 4). Figure 8 shows tone shading applied

together with colors from a transfer function. The tone effects are subtler, but still improve shape perception. The basic tissue colors are preserved, but the banded structure of the aorta is more apparent than in a simple illuminated and color-mapped image (Figure 2). Tone shading parameters used in Figures 7 and 8 are $k_{ty} = 0.3$, $k_{tb} = 0.3$, $k_{ta} = 1.0$, $k_{td} = 0.6$.

6 Application Examples

We have also applied the techniques in the previous sections to several other scientific data sets. Figures 10 and 11 are volume rendered images from a 256x256x64 MRI dataset of a tomato from Lawrence Berkeley National Laboratories. Figure 10 is a normal gas-based volume rendering of the tomato where a few of the internal structures are visible. Figure 11 has our volume illustration gradient and silhouette enhancements applied, resulting in a much more revealing image showing the internal structures within the tomato. Parameters used in Figure 11 are $k_{gc} = 0.5$, $k_{gs} = 2.5$, $k_{ge} = 3.0$; $k_{sc} = 0.4$, $k_{ss} = 500$, $k_{se} = 0.3$.

Figure 12 shows a 512x512x128 element flow data set from the time series simulation of unsteady flow emanating from a 2D rectangular slot jet. The 2D jet source is located at the left of the image and the flow is to the right. Flow researchers notice that both Figures 12 and 13 resemble Schlieren photographs that are traditionally used to analyze flow. Figure 13 shows the effectiveness of boundary enhancement, silhouette enhancement, and tone shading on this data set. The overall flow structure, vortex shedding, and helical structure are much easier to perceive in Figure 13 than in Figure 12.

Figures 14 and 15 are volume renderings of a 64x64x64 high-potential iron protein data set. Figure 14 is a traditional gas-based rendering of the data. Figure 15 has our tone shading volume illustration techniques applied, with parameters $k_{ty} = 0.15$, $k_{tb} = 0.15$, $k_{ta} = 1.0$, $k_{td} = 0.6$. The relationship of structure features and the three-dimensional location of the features is much clearer with the tone-based shading enhancements applied.

7 Conclusions

We have introduced the concept of volume illustration, combining the strengths of direct volume rendering with the expressive power of non-photorealistic rendering techniques. Volume illustration provides a powerful unified framework for producing a wide range of illustration styles using local and global properties of the volume model to control opacity accumulation and illumination. Volume illustration techniques enhance the perception of structure, shape, orientation, and depth relationships in a volume model. Comparing standard volume rendering (Figures 2, 10, 12, 14) with volume illustration images (Figures 3, 4, 5, 6, 7, 8, 9, 11, 13, 15) clearly shows the power of employing volumetric illustration techniques to enhance 3D depth perception and volumetric feature understanding.

8 Future Directions

We plan on extending our collection of NPR techniques and exploring suitability of these volume illustration techniques for data exploration and diagnosis.

9 Acknowledgements

We would like to thank researchers at the Mississippi State University NSF Computational Field Simulation Engineering Research Center and the Armed Forces Institute of Pathology for help in evaluating the effectiveness of these technique and

To appear in *Proceedings of IEEE Visualization '00* (October 2000, Salt Lake City, UT).

guiding our research. We would also like to thank Dr. Elliott Fishman of Johns Hopkins Medical Institutions for the abdominal CT dataset. The iron protein dataset came from the vtk website (www.kitware.com/vtk.html). Christopher Morris generated some of the pictures included in this paper. This work supported in part by the National Science Foundation under Grants ACIR 9996043 and ACIR 9978032.

References

- [Bierstadt1881] Albert Bierstadt. "Near Salt Lake City, Utah," Museum of Art, Brigham Young University, 1881.
- [Clark99] John O.E. Clark. *A Visual Guide to the Human Body, Barnes and Noble Books, 1999.*
- [daVinci1506] Leonardo daVinci, "The Virgin of the Rocks," National Gallery, London, 1503-1506.
- [Drebin88] Robert A. Drebin and Loren Carpenter and Pat Hanrahan. Volume Rendering, *Computer Graphics (Proceedings of SIGGRAPH 88)*, 22(4), pp. 65-74 (August 1988, Atlanta, Georgia). Edited by John Dill.
- [Ebert90] David S. Ebert and Richard E. Parent. Rendering and Animation of Gaseous Phenomena by Combining Fast Volume and Scanline A-buffer Techniques, *Computer Graphics (Proceedings of SIGGRAPH 90)*, 24 (4), pp. 357-366 (August 1990, Dallas, Texas). Edited by Forest Baskett. ISBN 0-201-50933-4.
- [Fang98] Shiao-fen Fang and Tom Biddlecome and Mihran Tuceryan. Image-Based Transfer Function Design for Data Exploration in Volume Visualization, *IEEE Visualization '98*, pp. 319-326 (October 1998). IEEE. Edited by David Ebert and Hans Hagen and Holly Rushmeier. ISBN 0-8186-9176-X.
- [Foley96] James Foley, Andries van Dam, Steven Feiner, and John Hughes, *Computer Graphics: Principles and Practice, Second Edition in C*, Addison Wesley 1996.
- [Fujishiro99] Issei Fujishiro and Taeko Azuma and Yuriko Takeshima. Automating Transfer Function Design for Comprehensible Volume Rendering Based on 3D Field Topology Analysis, *IEEE Visualization '99*, pp. 467-470 (October 1999, San Francisco, California). IEEE. Edited by David Ebert and Markus Gross and Bernd Hamann. ISBN 0-7803-5897-X.
- [Gooch98] 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, ACM SIGGRAPH, ACM Press, July 1998.
- [Gooch99] Bruce Gooch and Peter-Pike J. Sloan and Amy Gooch and Peter Shirley and Rich Riesenfeld. Interactive Technical Illustration, *1999 ACM Symposium on Interactive 3D Graphics*, pp. 31-38 (April 1999). ACM SIGGRAPH. Edited by Jessica Hodgins and James D. Foley. ISBN 1-58113-082-1.
- [Interrante95] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Enhancing Transparent Skin Surfaces with Ridge and Valley Lines, *IEEE Visualization '95*, pp. 52-59 (October 1995, Atlanta GA). IEEE. Edited by Gregory Nielson and Deborah Silver. ISBN 0-8186-7187-4.
- [Interrante97] Victoria Interrante and Henry Fuchs and Stephen M. Pizer. Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture, *IEEE Transactions on Visualization and Computer Graphics*, 3(2), (April - June 1997). ISSN 1077-2626.
- [Interrante98] Victoria Interrante and Chester Grosch. Visualizing 3D Flow, *IEEE Computer Graphics & Applications*, 18(4), pp. 49-53 (July - August 1998). ISSN 0272-1716.
- [Kajiya84] James T. Kajiya and Brian P. Von Herzen. Ray Tracing Volume Densities, *Computer Graphics (Proceedings of SIGGRAPH 84)*, 18(3), pp. 165-174 (July 1984, Minneapolis, Minnesota). Edited by Hank Christiansen.
- [Kindlmann98] Gordon Kindlmann and James Durkin. Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering, In *Proceedings of 1998 IEEE Symposium on Volume Visualization*, pp. 79-86.
- [Kirby99] 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 (October 1999, San Francisco, California). IEEE. Edited by David Ebert and Markus Gross and Bernd Hamann. ISBN 0-7803-5897-X.
- [Krueger91] Wolfgang Krueger, The Application of transport theory to the visualization of 3D scalar fields, *Computers in Physics*, pp. 397-406, July 1991.
- [Laidlaw98] David H. Laidlaw and Eric T. Ahrens and David Kremers and Matthew J. Avalos and Russell E. Jacobs and Carol Readhead. Visualizing Diffusion Tensor Images of the Mouse Spinal Cord, *IEEE Visualization '98*, pp. 127-134 (October 1998). IEEE. Edited by David Ebert and Hans Hagen and Holly Rushmeier. ISBN 0-8186-9176-X.
- [Levoy90] Marc Levoy. Efficient Ray Tracing of Volume Data, *ACM Transactions on Graphics*, 9 (3), pp. 245-261 (July 1990). ISSN 0730-0301.
- [Max95] Nelson Max. Optical models for direct volume rendering, *IEEE Transactions on Visualization and Computer Graphics*, 1 (2), pp. 99-108 (June 1995). ISSN 1077-2626.
- [Nishita87] Tomoyuki Nishita and Yasuhiro Miyawaki and Eihachiro Nakamae. A Shading Model for Atmospheric Scattering Considering Luminous Intensity Distribution of Light Sources, *Computer Graphics (Proceedings of SIGGRAPH 87)*, 21 (4), pp. 303-310 (July 1987, Anaheim, California). Edited by Maureen C. Stone.
- [Nishita98] Tomoyuki Nishita. Light Scattering Models for the Realistic Rendering of Natural Scenes, *Eurographics Rendering Workshop 1998*, pp. 1-10 (June 1998, Vienna, Austria). Eurographics. Edited by George Drettakis and Nelson Max. ISBN3-211-83213-0.
- [Rheingans96] Penny Rheingans. Opacity-modulating Triangular Textures for Irregular Surfaces, *Proceedings of IEEE Visualization '96*, pp. 219-225 (October 1996, San Francisco CA). IEEE. Edited by Roni Yagel and Gregory Nielson. ISBN 0-89791-864-9.
- [Saito90] Takafumi Saito and Tokiichiro Takahashi. Comprehensible Rendering of 3-D Shapes, *Computer Graphics (Proceedings of SIGGRAPH 90)*, 24 (4), pp. 197-206 (August 1990, Dallas, Texas).
- [Saito94] Takafumi Saito. Real-time Previewing for Volume Visualization. In *Proceedings of 1994 IEEE Symposium on Volume Visualization*, pp. 99-106.
- [Salisbury94] Michael P. Salisbury and Sean E. Anderson and Ronen Barzel and David H. Salesin. Interactive Pen-And-Ink Illustration, *Proceedings of SIGGRAPH 94, Computer Graphics Proceedings, Annual Conference Series*, pp. 101-108 (July 1994, Orlando, Florida). ACM Press. Edited by Andrew Glassner. ISBN 0-89791-667-0.
- [Salisbury97] Michael P. Salisbury and Michael T. Wong and John F. Hughes and David H. Salesin. Orientable Textures for Image-Based Pen-and-Ink Illustration, *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual

To appear in *Proceedings of IEEE Visualization '00* (October 2000, Salt Lake City, UT).

Conference Series, pp. 401-406 (August 1997, Los Angeles, California). Addison Wesley. Edited by Turner Whitted. ISBN 0-89791-896-7.

[Treavett00] S.M.F. Treavett and M. Chen. Pen-and-Ink Rendering in Volume Visualisation, *Proceedings of IEEE Visualization 2000*, October 2000, ACM SIGGRAPH Press.

[Williams92] Peter L. Williams and Nelson Max. A Volume Density Optical Model, *1992 Workshop on Volume Visualization*, pp. 61-68 (1992). ACM.

[Winkenbach94] Georges Winkenbach and David H. Salesin. Computer-Generated Pen-And-Ink Illustration, *Proceedings of SIGGRAPH 94, Computer Graphics Proceedings, Annual Conference Series*, pp. 91-100 (July 1994, Orlando, Florida). ACM Press. Edited by Andrew Glassner. ISBN 0-89791-667-0.



Figure 1. Gaseous illumination of medical CT volume. Voxels are a constant color.




Figure 5. Volumetric sketch lines on CT volume. Lines are all white.



Figure 12. Atmospheric volume rendering of square jet. No illustration enhancements.




Figure 2. Gaseous illumination of color-mapped CT volume.




Figure 3. Color-mapped gaseous illumination with boundary enhancement.




Figure 4. Silhouette and boundary enhancement of CT volume.



Figure 6. Distance color blending and halos around features of CT volume.




Figure 7. Tone shading in CT volume. Surfaces toward light receive warm color.




Figure 8. Tone shading in colored volume. Surfaces toward light receive warm color.




Figure 9. Orientation fading. Surfaces toward viewer are desaturated.




Figure 10. Standard atmospheric volume rendering of tomato.




Figure 11. Boundary and silhouette enhanced tomato.




Figure 13. Square jet with boundary and silhouette enhancement, and tone shading.




Figure 14. Atmospheric rendering of iron protein.




Figure 15. Tone shaded iron protein.