Volume Illustration: Nonphotorealistic Rendering of Volume Models

Penny Rheingans, *Member*, *IEEE Computer Society*, and
David Ebert, *Member*, *IEEE Computer Society*

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 nonphotorealistic 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.

Index Terms—Volume rendering, nonphotorealistic rendering, illustration, lighting models, shading, transfer functions, visualization.

✦

---

1 INTRODUCTION

$\mathbf{F}$OR 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 [17], [5]. 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 [23]. 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 nonphotorealistic rendering (NPR) techniques [38], [34] 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. Nonphotorealistic 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.

- *P. Rheingans is with the Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, Baltimore, MD 21250. E-mail: rheingan@cs.umbc.edu.*
- *D. Ebert is with the Electrical and Computer Engineering Department, Purdue University, West Lafayette, IN 47907-1285. E-mail: ebertd@purdue.edu.*

2   RELATED WORK

Traditional volume rendering spans a spectrum from the accurate to the ad hoc. Kajiya and von Herzen's original work on volume ray tracing for generating images of clouds [17] incorporated a physics-based illumination and atmospheric attenuation model. This work in realistic volume rendering techniques has been extended by numerous researchers [29], [6], [20], [37], [25], [30]. In contrast, traditional volume rendering has relied on the use of transfer functions from scalar value to rendered opacity to produce artificial views of the data which highlight regions of interest [5]. 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 [7], [9]. 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. A few researchers have investigated the potential for using gradient information along with value in the transfer function, starting with Levoy's use of gradient magnitude to increase opacity in boundary regions [22]. Kindlmann and Durkin used the first and second directional derivatives in the gradient direction to calculate a boundary emphasis to be included in the opacity transfer function [18].

Additionally, there has been extensive research for illustrating surface shape using nonphotorealistic 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 [11]. The extraction and rendering of silhouettes and other expressive lines has been addressed by several researchers (including [32], [34], [12], [14], [10]). Expressive textures have also been applied to surfaces to convey surface shape (including [23], [31], [35], [15]).

A few researchers have applied NPR techniques to the display of data. Laidlaw et al. used concepts from painting to create visualizations of 2D data, using brushstroke-like elements to convey information [21] and a painterly process to compose complex visualizations [19]. Treavett and Chen have developed techniques for pen-and-ink illustrations of surfaces within volumes [36]. Interrante and Grosch applied principles from technical illustration to convey depth relationships with halos around foreground features in flow data [16]. Saito converted 3D scalar fields into a sampled point representation and visualized selected points with a simple primitive, creating an NPR look [33]. However, with few exceptions, 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., [23]). In contrast, our
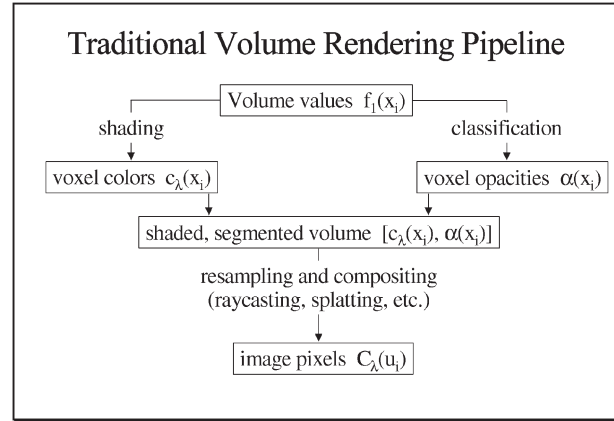

Diagram 1. Traditional volume rendering pipeline.

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 allow 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 illustrative effects.
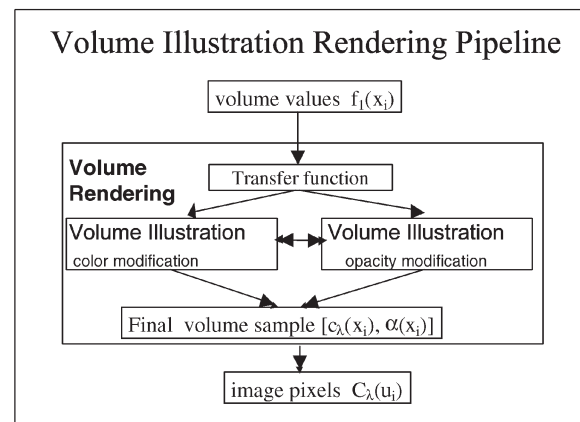

Diagram 2. Volume illustration rendering pipeline.

TABLE 1
Types of Enhancements

Feature Enhancement	Depth and Orientation Cues	Regional Enhancement
• Boundary Enhancement • Silhouettes • Fading • Sketch Lines	• Distance Color Blending • Feature Halos • Tone Shading	Regional Enhancement

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 and includes volumetric shadowing and self-shadowing [6]. Fig. 3 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 Fig. 4, 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 to better show detail within semitransparent volumes, but the same approach can be easily applied to a base renderer using Phong illumination and linear accumulation.

In the following sections, we describe our current collection of volume illumination 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 propose are of three basic types:

- feature enhancement,
- depth and orientation cues,
- regional enhancement.

The specific techniques explored are listed in Table 1.

4 FEATURES

Assume a scalar 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 $\nabla_f (P_i)$ at that location. In many operations, we will want that gradient to be normalized. We use ∇_{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. 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. If no transfer function is used, these values can be set to constants for the whole volume. The inclusion of a transfer function allows illustrative enhancements to supplement, rather than replace, existing volume rendering mechanisms. For the images in this paper, we have chosen to calculate the gradient used for enhancement from the datasets after the simple power function transfer function has been applied. Our experimentation with these datasets has shown that this gradient is more effective in capturing subtle features within the volume.

Numerous researchers have discussed the importance of gradient calculation methods to high quality volume rendering results. The central difference method for estimating voxel gradients is simple and commonly used. In this method, the voxel values along each dimension are compared to produce the gradient. Specifically,

$$\nabla_f(x_i, y_i, z_i) = \begin{bmatrix} \frac{1}{2}f(x_{i-1}, y_i, z_i) - \frac{1}{2}f(x_{i+1}, y_i, z_i) \\ \frac{1}{2}f(x_i, y_{i-1}, z_i) - \frac{1}{2}f(x_i, y_{i+1}, z_i) \\ \frac{1}{2}f(x_i, y_i, z_{i-1}) - \frac{1}{2}f(x_i, y_i, z_{i+1}) \end{bmatrix}.$$

Such a filter is linearly separable in that the calculation of each gradient component depends only on the voxel values along that component dimension. Several researchers have analyzed the characteristics of linearly separable gradient filters (for example, [13], [26]). Möller et al. proposed a methodology for designing linearly separable derivative filters with specific accuracy and continuity characteristics using optimized piecewise polynomial interpolation filters [27]. Naturally, this method has accuracy and smoothness advantages for estimating gradients both at volume sample points and within the voxel, but here we will consider only the results at volume sample points. Using an error function of second degree (2EF), the filter reduces to the central difference since the higher order terms producing the required continuity only contribute to the result at locations within the voxel. Using 3EF constraints produces the following filter for gradients at the sample points:
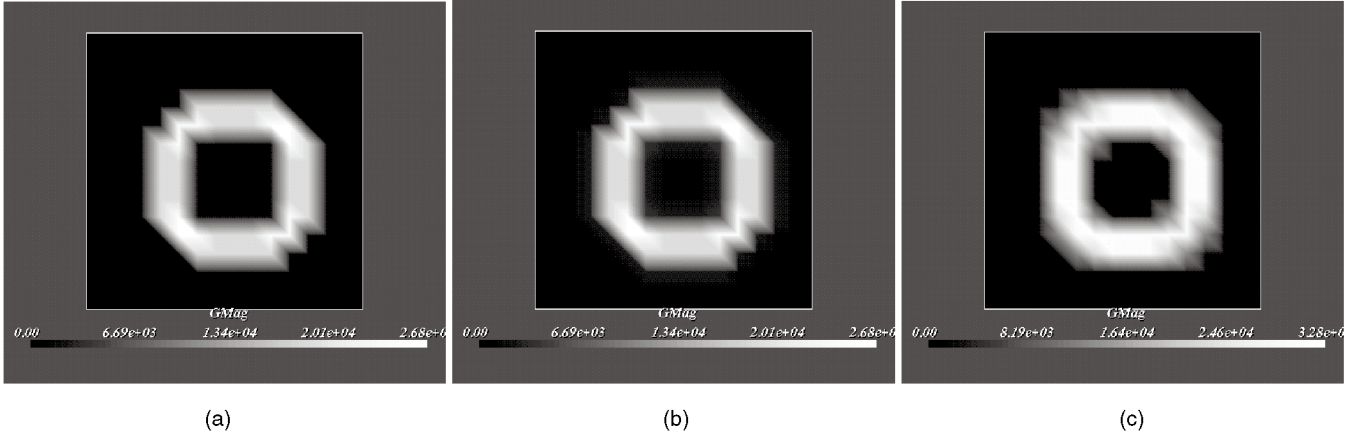
Fig. 1. Gradient magnitude slice estimated using (a) central differences, (b) linearly separable filter, (c) regression.

$$\nabla_f(x_i, y_i, z_i) =$$
$$\begin{bmatrix} -\frac{1}{12}f(x_{i-2}, y_i, z_i) + \frac{2}{3}f(x_{i-1}, y_i, z_i) - \\ \frac{2}{3}f(x_{i+1}, y_i, z_i) + \frac{1}{12}f(x_{i+2}y_i, z_i) \\ -\frac{1}{12}f(x_i, y_{i-2}, z_i) + \frac{2}{3}f(x_i, y_{i-1}, z_i) - \\ \frac{2}{3}f(x_i, y_{i+1}, z_i) + \frac{1}{12}f(x_i, y_{i+2}, z_i) \\ -\frac{1}{12}f(x_i, y_i, z_{i-2}) + \frac{2}{3}f(x_i, y_i, z_{i-1}) - \\ \frac{2}{3}f(x_i, y_i, z_{i+1}) + \frac{1}{12}f(x_i, y_i, z_{i+2}) \end{bmatrix}.$$

Other research has analyzed the capabilities of both separable and the more computationally expensive non-separable filters (for example, [24], [2]). Neumann et al. propose a method they call 4D linear regression for gradient estimation, which approximates the density function in a local neighborhood with a 3D regression hyperplane [28]. Plane equation coefficients (and, hence, gradient components) are calculated so as to minimize the error from the 26 neighboring voxels weighted by a spherically symmetric function. This yields plane equations of:

$$A = w_A \sum_{k=0}^{26} w_k v_k x_k$$

$$B = w_B \sum_{k=0}^{26} w_k v_k y_k$$

$$C = w_C \sum_{k=0}^{26} w_k v_k z_k$$

$$D = w_D \sum_{k=0}^{26} w_k v_k,$$

where:

$$w_A = \frac{1}{\sum_{k=0}^{26} w_k x_k^2}, w_B = \frac{1}{\sum_{k=0}^{26} w_k y_k^2},$$

$$w_C = \frac{1}{\sum_{k=0}^{26} w_k z_k^2}, w_D = \frac{1}{\sum_{k=0}^{26} w_k},$$

and the weighting function is inversely proportional to the Euclidean distance between the voxel and the neighbor:

$$w_k = \frac{1}{d_k}.$$

The voxel gradient is then:

$$\nabla_f = \begin{bmatrix} A \\ B \\ C \end{bmatrix}.$$

We have experimented with these three different gradient estimation methods. Fig. 1a shows the gradient magnitude through a representative slice of a $16 \times 16 \times 16$ spherical volume calculated using central differences. Fig. 1b shows the gradient magnitude through the same slice calculated using linearly separable piecewise polynomial interpolation (from Möller et al.). Fig. 1c shows the gradient magnitude calculated using 4D linear regression (from Neumann et al.). The gradients estimated using central differences and polynomial interpolation are very similar, but not identical. Gradients calculated by polynomial interpolation have slightly softer edges in regions of rapid change. Gradients calculated using linear regression show more smoothing in regions where gradient direction is changing rapidly. Fig. 2 shows the results of standard gradient-based volume rendering of this volume using gradients estimated by the three methods. The images computed using central differences (Fig. 2a) and polynomial interpolation (Fig. 2b) are similar, displaying substantial voxel artifact from the very coarse data resolution. The image computed using regression (Fig. 2c) shows a much smoother appearance. Our initial observations on more complex data sets indicate that, while the volume illustration images created using different gradient methods do differ, the differences do not result in a systematic difference in expressive power. Further analysis of the effects of gradient estimation method for volume illustration is a topic for future consideration. All figures beginning with Fig. 3 use central differences to estimate gradient.

5 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 "surfaceness" a Boolean quality. Many other features, such as silhouettes or regions of high curvature,

(a)   (b)   (c)
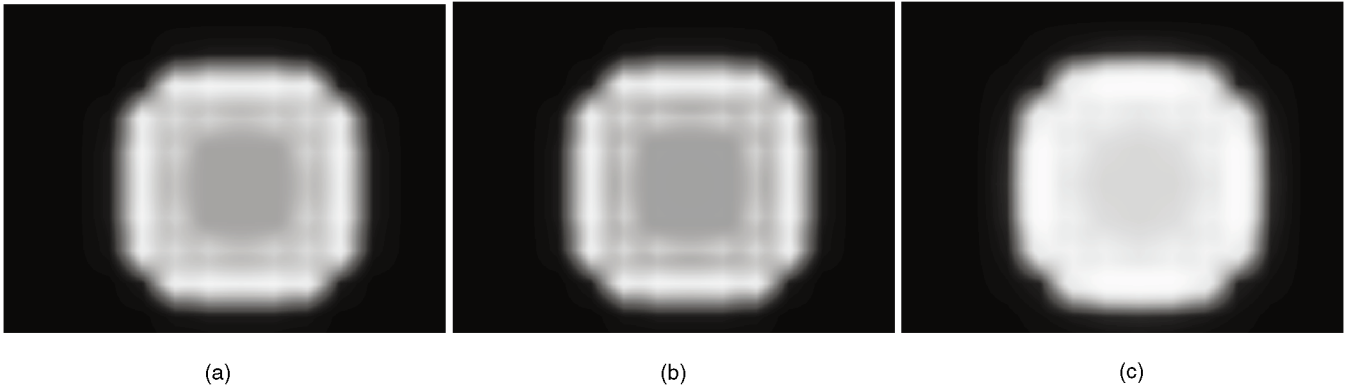
Fig. 2. Sphere rendering using (a) central differences, (b) linearly separable filter, (c) regression.

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 "surfaceness" 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.

5.1 Boundary Enhancement

Levoy [22] 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.

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:

$$o_g = o_v \left(k_{gc} + k_{gs} \left(\|\nabla_{fn}\| \right)^{k_{gc}} \right),$$

where o_v is the original opacity and ∇_{fn} 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 ($k_{gs} \geq 1$) 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.

Fig. 5 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 Fig. 5 are $k_{gc} = 0.7$, $k_{gs} = 10$, $k_{ge} = 2.0$.
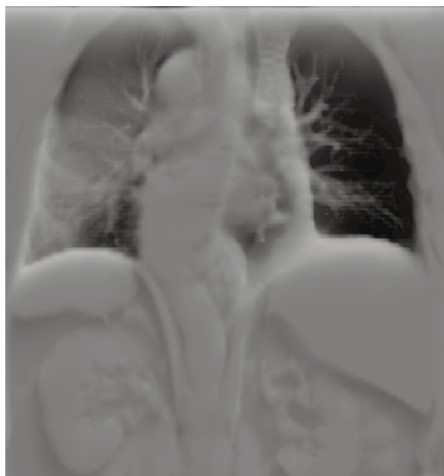

Fig. 3. Gaseous illumination of medical CT volume. Voxels are a constant color.

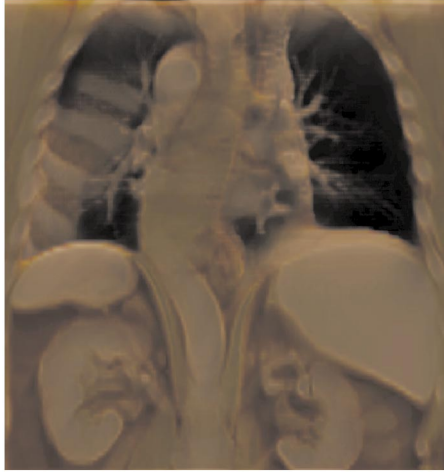Fig. 4. Gaseous illumination of color-mapped CT volume.

Fig. 5. Color-mapped gaseous illumination with boundary enhancement.

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

Surface orientation is an important visual cue that has been successfully conveyed by illustrators 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 [34], [14]. 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:

$$o_s = o_v\Big(k_{sc} + k_{ss}\big(1 - abs(\nabla_{fn} \cdot V)\big)^{k_{se}}\Big),$$

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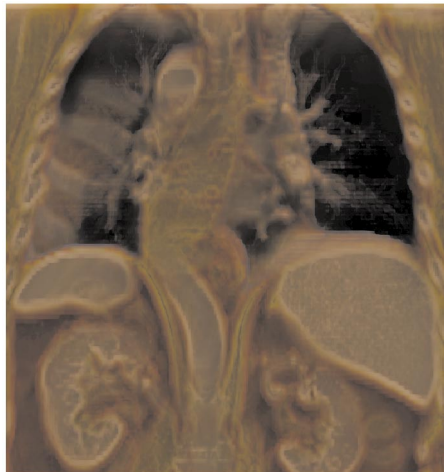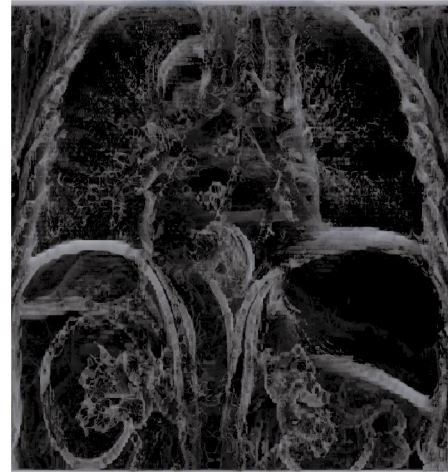

Fig. 7. Volumetric sketch lines on CT volume. Lines are all white.

Fig. 6 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 Fig. 6 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 Fig. 7. Fig. 7 shows a black and white sketch of the medical dataset by using a white sketch color and making nonvolumetric 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. Fig. 8 shows a colored sketch.

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 [3]. We simulate these effects by allowing the volumetric gradient orientation to

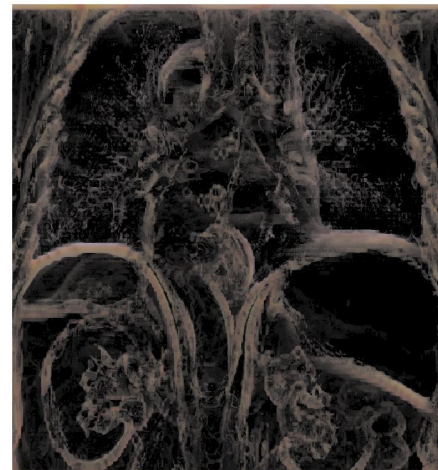Fig. 6. Silhouette and boundary enhancement of CT volume.

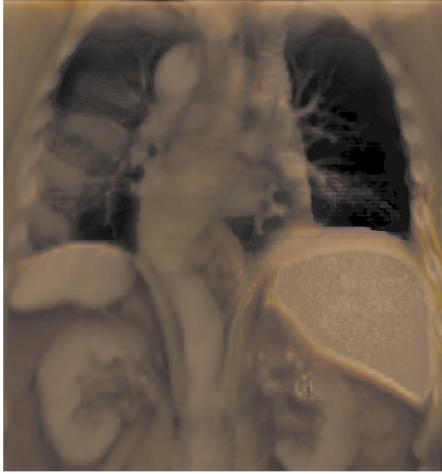Fig. 8. Volumetric sketch lines on CT volume. Lines colored by colormap.

Fig. 9. Regional enhancement centered on liver.

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. Fig. 14 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.

6 REGIONAL ENHANCEMENT

Technical illustrators selectively add detail and enhancement to emphasize areas within their illustrations. We have adapted this idea to volume illustration by performing enhancements only within a specified region of a dataset (e.g., an organ of interest). Regional enhancement can be easily implemented by calculating an enhancement scalar that determines the amount of each enhancement technique that is applied during rendering.

We have chosen to allow the user to specify three simple parameters to apply regional enhancement:

1. enhancement center of interest (ECOI),
2. enhancement maximum (k_{rs}),
3. enhancement fall-off power (k_{re}),

The parameters are used in the following equation to determine the enhancement amount, e:

$$e = k_{rs} * (\nabla_{fn} \cdot E)^{k_{re}},$$

where $E = (\text{eye} - \text{ECOI})/\|(\text{eye} - \text{ECOI})\|$. By choosing small values of k_{re} (< 2), the regional enhancement decreases gradually from the center of enhancement while larger values of k_{re} (e.g., 20, 50) produce sharp, focused enhanced regions.

The regional enhancement is very effective in immediately drawing the user's attention to important regions within an image. Fig. 9 shows an enhanced abdomen CT rendering, with the enhancement region centered on the liver. Fig. 10 shows a volume rendering of a segmented kidney CT dataset with the center of regional enhancement set to the center of the right kidney. The internal structures of the right kidney are clearly apparent.

7 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 [6], [17]. Accurate volumetric shadowing often produces subtle effects which do not provide strong three-dimensional depth cues. As a result, the shape of individual

Fig. 10. Regional enhancement of right kidney.

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.

7.1 Distance Color Blending

Intensity depth-cueing is a well-known technique for enhancing the perception of depth in a scene [8]. 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-cueing 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 [4], [1]. 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-cueing need not be linear; we use an exponential function to control the application of depth-cueing. The distance color blending process can be described by:

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

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 gray ($c_b = (a, a, a)$ for some value of a) only standard intensity depth-cueing 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.

Fig. 11 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 Fig. 11 are $c_b = (0, 0, 0.15)$, $k_{ds} = 1.0$, $k_{se} = 0.5$.

7.2 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 and Grosch [16] used a similar idea to show depth relationships in 3D flow data using Line Integral Convolution (LIC). They 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.

Fig. 11. Distance color blending and halos around features of CT volume.

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:

$$h(i) = \left(\sum_{n}^{neighbors} \frac{h_n}{\|P_i - P_n\|^2}\right)\left(1 - \|\nabla_f(P_i)\|\right),$$

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:

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

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 strongest 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 and Grosch, 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.

Fig. 11 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 Fig. 11 are $k_{hpe} = 1.0$ and $k_{hse} = 2.0$.

7.3 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 [11], 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
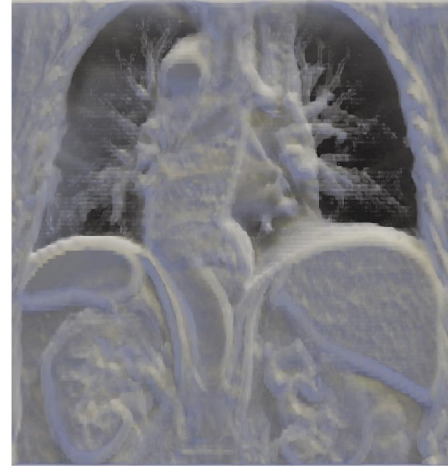

Fig. 12. Tone shading in CT volume. Surfaces toward light receive warm color.

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:

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

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

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

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.

Fig. 12 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 (Fig. 6). Fig. 13 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 (Fig. 2). Tone shading parameters used in Figs. 12 and 13 are $k_{ty} = 0.3$, $k_{tb} = 0.3$, $k_{ta} = 1.0$, $k_{td} = 0.6$.
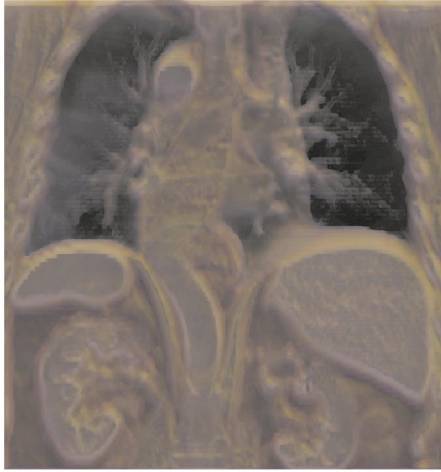
Fig. 13. Tone shading in colored volume. Surfaces toward light receive warm color.

8   APPLICATION EXAMPLES

We have also applied the techniques in the previous sections to several other scientific data sets. Figs. 15 and 16 are volume rendered images from a $256 \times 256 \times 64$ MRI dataset of a tomato from Lawrence Berkeley National Laboratories. Fig. 15 is a normal gas-based volume rendering of the tomato where a few of the internal structures are visible. Fig. 16 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 Fig. 16 are $k_{gc} = 0.5$, $k_{gs} = 2.5$, $k_{ge} = 3.0$; $k_{sc} = 0.4$, $k_{ss} = 500$, $k_{se} = 0.3$.

Fig. 17 shows a $512 \times 512 \times 128$ 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 Figs. 17 and 18 resemble Schlieren photographs that are traditionally used to analyze flow. Fig. 18 shows the effectiveness of boundary enhancement, silhouette enhancement, and tone shading on this data set. The overall flow structure, vortex shedding, and
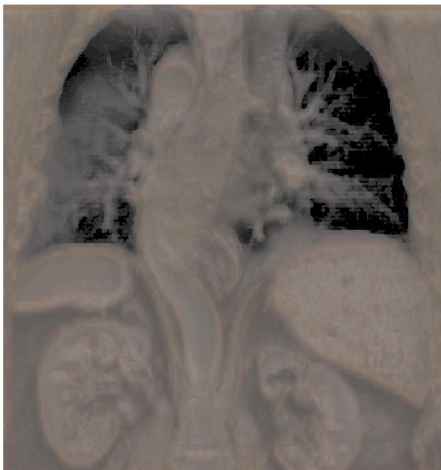

Fig. 15. Standard atmospheric volume rendering of a tomato.

helical structure are much easier to perceive in Fig. 18 than in Fig. 17.

Figs. 19 and 20 are volume renderings of a $64 \times 64 \times 64$ high-potential iron protein data set. Fig. 14 is a traditional gas-based rendering of the data. Fig. 20 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.

9   DISCUSSION

To some extent, volume illustration techniques may be regarded as representing a very complex and flexible multivariate transfer function mechanism. Using this abstraction, the parameters controlling boundary and silhouette enhancement specify a multivariate opacity transfer function from voxel value, voxel location, gradient, and view vector. Similarly, the parameters controlling oriented fading, distance color blending, and tone shading specify a multivariate color transfer function from voxel value, initial voxel color, voxel location, gradient, view vector, and light position and intensity. The two transfer functions would require 10 and 17 inputs, respectively. Specifying illustration techniques through multiple separable enhancements breaks the problem of designing these potentially very complex multivariate functions into smaller, more manageable design problems.

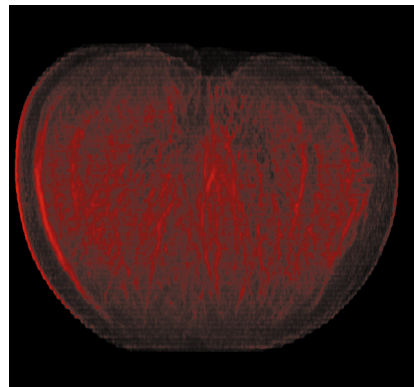Fig. 14. Orientation fading. Surfaces toward viewer are desaturated.

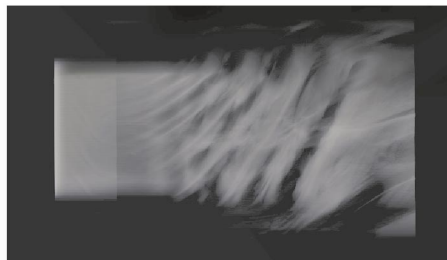Fig. 16. Boundary and silhouette enhanced tomato.

Fig. 17. Atmospheric volume rendering of a square jet. No illustration enhancements.

The transfer function abstraction breaks down when feature halos are considered. Both transfer functions would not only need to take the voxel value, voxel location, gradient, and view vector as inputs, but also the value, location, and gradient of each voxel in a neighborhood of user-controllable size. This would require transfer functions with a number of inputs which was not only large, but variable.

10 CONCLUSIONS

We have introduced the concept of volume illustration, combining the strengths of direct volume rendering with the expressive power of nonphotorealistic 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 (Figs. 4, 15, 17, 19) with volume illustration images (Figs. 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 16, 18, 20) clearly shows the power of employing volumetric illustration techniques to enhance 3D depth perception and volumetric feature understanding.

11 FUTURE DIRECTIONS

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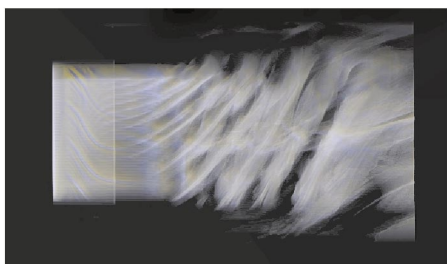

Fig. 19. Atmospheric rendering of iron protein.

REFERENCES

[1] A. Bierstadt, "Near Salt Lake City, Utah," Museum of Art, Brigham Young Univ., 1881.
[2] M. Bentum, B. Lichtenbelt, and T. Malzbender, "Frequency Analysis of Gradient Estimators in Volume Rendering," *IEEE Trans. Visualization and Computer Graphics,* vol. 2, no. 3, Sept. 1996.
[3] J.O.E. Clark, *A Visual Guide to the Human Body.* Barnes and Noble Books, 1999.
[4] L. daVinci, "The Virgin of the Rocks," National Gallery, London, 1503-1506.
[5] R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics (Proc. SIGGRAPH 88),* J. Dill, ed., vol. 22, no. 4, pp. 65-74, Aug. 1988.
[6] D.S. Ebert and R.E. Parent, "Rendering and Animation of Gaseous Phenomena by Combining Fast Volume and Scanline A-Buffer
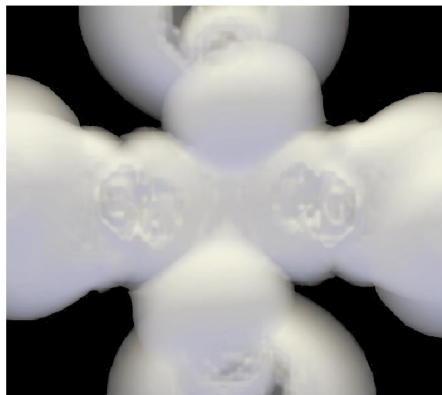
Fig. 18. Square jet with boundary and silhouette enhancement and tone shading.

Fig. 20. Tone shaded iron protein.

Techniques," *Computer Graphics (Proc. SIGGRAPH 90),* F. Baskett, ed., vol. 24, no. 4, pp. 357-366, Aug. 1990.

[7] S. Fang, T. Biddlecome, and M. Tuceryan, "Image-Based Transfer Function Design for Data Exploration in Volume Visualization," *Proc. IEEE Visualization '98,* D. Ebert, H. Hagen, and H. Rushmeier, eds., pp. 319-326, Oct. 1998.

[8] J. Foley, A. van Dam, S. Feiner, and J. Hughes, *Computer Graphics: Principles and Practice, Second Edition in C,* Addison-Wesley, 1996.

[9] I. Fujishiro, T. Azuma, and Y. Takeshima, "Automating Transfer Function Design for Comprehensible Volume Rendering Based on 3D Field Topology Analysis," *Proc. IEEE Visualization '99,* D. Ebert, M. Gross, and B. Hamann, eds., pp. 467-470, Oct. 1999.

[10] A. Girshik, V. Interrante, S. Haker, and T. Lemoine, "Line Direction Matters: An Argument for the Use of Principal Curvature Directions in 3D Line Drawings," *Proc. First Int'l Symp. Non Photorealistic Animation and Rendering (NPAR2000),* pp. 43-52, 2000.

[11] A. Gooch, B. Gooch, P. Shirley, and E. Cohen, "A Non-Photorealistic Lighting Model for Automatic Technical Illustration," *Proc. SIGGRAPH '98, Computer Graphics Proc., Ann. Conf. Series,* pp. 447-452, July 1998.

[12] B. Gooch, P.-P.J. Sloan, A. Gooch, P. Shirley, and R. Riesenfeld, "Interactive Technical Illustration," *Proc. 1999 ACM Symp. Interactive 3D Graphics,* J. Hodgins and J.D. Foley, eds., pp. 31-38, Apr. 1999.

[13] M. Goss, "An Adjustable Gradient Filter for Volume Visualization Image Enhancement," *Proc. Graphics Interface '94,* pp. 67-74, 1994.

[14] V. Interrante, H. Fuchs, and S. Pizer, "Enhancing Transparent Skin Surfaces with Ridge and Valley Lines," *Proc. IEEE Visualization '95,* G. Nielson and D. Silver, eds, pp. 52-59, Oct. 1995.

[15] V. Interrante, H. Fuchs, and S.M. Pizer, "Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture," *IEEE Trans. Visualization and Computer Graphics,* vol. 3, no. 2, Apr.-June 1997.

[16] V. Interrante and C. Grosch, "Visualizing 3D Flow," *IEEE Computer Graphics and Applications,* vol. 18, no. 4, pp. 49-53, July-Aug. 1998.

[17] J.T. Kajiya and B.P. Von Herzen, "Ray Tracing Volume Densities," *Computer Graphics (Proc. SIGGRAPH '84),* H. Christiansen, ed., vol. 18, no. 3, pp. 165-174, July 1984.

[18] G. Kindlmann and J. Durkin, "Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering," *Proc. 1998 IEEE Symp. Volume Visualization,* pp. 79-86, 1998.

[19] R.M. Kirby, H. Marmanis, and D.H. Laidlaw, "Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting," *Proc. IEEE Visualization '99,* D. Ebert, M. Gross, and B. Hamann, eds., pp. 333-340, Oct. 1999.

[20] W. Krueger, "The Application of Transport Theory to the Visualization of 3D Scalar Fields," *Computers in Physics,* pp. 397-406, July 1991.

[21] D.H. Laidlaw, E.T. Ahrens, D. Kremers, M.J. Avalos, R.E. Jacobs, and C. Readhead, "Visualizing Diffusion Tensor Images of the Mouse Spinal Cord," *Proc. IEEE Visualization '98,* D. Ebert, H. Hagen, and H. Rushmeier, eds., pp. 127-134, Oct. 1998.

[22] M. Levoy, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications,* vol. 8, no. 5, pp. 29-37, 1988.

[23] M. Levoy, "Efficient Ray Tracing of Volume Data," *ACM Trans. Graphics,* vol. 9, no. 3, pp. 245-261, July 1990.

[24] S. Marschner and R. Lobb, "An Evaluation of Reconstruction Filters for Volume Rendering," *Proc. IEEE Visualization '94,* pp. 100-107, 1994.

[25] N. Max, "Optical Models for Direct Volume Rendering," *IEEE Trans. Visualization and Computer Graphics,* vol. 1, no. 2, pp. 99-108, June 1995.

[26] T. Möller, R. Machiraju, K. Mueller, and R. Yagel, "Evaluation and Design of Filters Using a Taylor Series Expansion," *IEEE Trans. Visualization and Computer Graphics,* vol. 3, no. 2, pp. 184-199, Apr.-June 1997.

[27] T. Möller, K. Mueller, Y. Kurzion, R. Machiraju, and R. Yagel, "Design of Accurate and Smooth Filters for Function and Derivative Reconstruction," *Proc. 1998 IEEE Symp. Volume Visualization,* pp. 143-151, 1998.

[28] L. Neumann, B. Csébfalvi, A. König, and E. Gröller, "Gradient Estimation in Volume Data Using 4D Linear Regression," *Proc. Eurographics 2000,* 2000.

[29] T. Nishita, Y. Miyawaki, and E. Nakamae, "A Shading Model for Atmospheric Scattering Considering Luminous Intensity Distribution of Light Sources," *Computer Graphics (Proc. SIGGRAPH '87),* M.C. Stone, ed., vol. 21, no. 4, pp. 303-310, July 1987.

[30] T. Nishita, "Light Scattering Models for the Realistic Rendering of Natural Scenes," *Proc. Eurographics Rendering Workshop,* G. Drettakis and N. Max, eds., pp. 1-10, June 1998.

[31] P. Rheingans, "Opacity-Modulating Triangular Textures for Irregular Surfaces," *Proc. IEEE Visualization '96,* R. Yagel and G. Nielson, eds., pp. 219-225, Oct. 1996.

[32] T. Saito and T. Takahashi, "Comprehensible Rendering of 3-D Shapes," *Computer Graphics (Proc. SIGGRAPH '90),* vol. 24, no. 4, pp. 197-206, Aug. 1990.

[33] T. Saito, "Real-Time Previewing for Volume Visualization," *Proc. 1994 IEEE Symp. Volume Visualization,* pp. 99-106, 1994.

[34] M.P. Salisbury, S.E. Anderson, R. Barzel, and D.H. Salesin, "Interactive Pen-and-Ink Illustration," *Proc. SIGGRAPH '94, Computer Graphics Proc., Ann. Conf. Series,* A. Glassner, ed., pp. 101-108, July 1994.

[35] M.P. Salisbury, M.T. Wong, J.F. Hughes, and D.H. Salesin, "Orientable Textures for Image-Based Pen-and-Ink Illustration," *Proc. SIGGRAPH '97, Computer Graphics Proc., Ann. Conf. Series,* T. Whitted, ed., pp. 401-406, Aug. 1997.

[36] S.M.F. Treavett and M. Chen, "Pen-and-Ink Rendering in Volume Visualisation," *Proc. IEEE Visualization 2000,* pp. 203-209, Oct. 2000.

[37] P.L. Williams and N. Max, "A Volume Density Optical Model," *Proc. 1992 Workshop Volume Visualization,* pp. 61-68, 1992.

[38] G. Winkenbach and D.H. Salesin, "Computer-Generated Pen-and-Ink Illustration," *Proc. SIGGRAPH '94, Computer Graphics Proc., Ann. Conf. Series,* A. Glassner, ed., pp. 91-100, July 1994.

Penny Rheingans received the PhD degree in computer science from the University of North Carolina, Chapel Hill, and the BS degree in computer science from Harvard University. She is an assistant professor of computer science at the University of Maryland Baltimore County. Her current research interests include uncertainty in visualization, multivariate visualization, volume visualization, information visualization, perceptual and illustration issues in visualization, dynamic and interactive representations and interfaces, and the experimental validation of visualization techniques. She is a member of the IEEE Computer Society.

David Ebert received the PhD degree from the Computer and Information Science Department at the Ohio State University in 1991. He is an associate professor in the School of Electrical and Computer Engineering at Purdue University. His research interests are scientific, medical, and information visualization, computer graphics, animation, and procedural techniques. Dr. Ebert performs research in volume rendering, nonphotorealistic visualization, minimally immersive visualization, realistic rendering, procedural texturing, modeling, and animation, modeling natural phenomena, and volumetric display software. He has also been very active in the graphics community, teaching courses, presenting papers, chairing the ACM SIGGRAPH '97 Sketches program, and cochairing the IEEE Visualization '98 and '99 papers program. He is a member of the IEEE Computer Society.

▷ **For further information on this or any computing topic, please visit our Digital Library at** http://computer.org/publications/dlib.