

Designing Effective Transfer Functions for Volume Rendering from Photographic Volumes

David S. Ebert, *Member, IEEE*, Christopher J. Morris,
Penny Rheingans, *Member, IEEE Computer Society*, and
Terry S. Yoo, *Member, IEEE Computer Society*

Abstract—Photographic volumes present a unique, interesting challenge for volume rendering. In photographic volumes, voxel color is predetermined, making color selection through transfer functions unnecessary. However, photographic data does not contain a clear mapping from the multivalued color values to a scalar density or opacity, making projection and compositing much more difficult than with traditional volumes. Moreover, because of the nonlinear nature of color spaces, there is no meaningful norm for the multivalued voxels. Thus, the individual color channels of photographic data must be treated as incomparable data tuples rather than as vector values. Traditional differential geometric tools, such as intensity gradients, density, and Laplacians, are distorted by the nonlinear nonorthonormal color spaces that are the domain of the voxel values. We have developed different techniques for managing these issues while directly rendering volumes from photographic data. We present and justify the normalization of color values by mapping RGB values to the CIE $L^*u^*v^*$ color space. We explore and compare different opacity transfer functions that map three channel color values to opacity. We apply these many-to-one mappings to the original RGB values as well as to the voxels after conversion to $L^*u^*v^*$ space. Direct rendering using transfer functions allows us to explore photographic volumes without having to commit to an a priori segmentation that might mask fine variations of interest. We empirically compare the combined effects of each of the two color spaces with our opacity transfer functions using source data from the Visible Human Project.

Index Terms—Volume rendering, transfer functions, photographic data.

1 INTRODUCTION

RECENTLY, several photographic volume data sets have become available from projects such as the “Whole Frog Project” at Lawrence Berkeley National Laboratory [22], the Visible Embryo Project at the Armed Forces Institute of Pathology, and the Visible Human Project (VHP) at the National Library of Medicine [27]. This type of data offers exciting possibilities for realistic volume visualization since correct color values are known for each voxel. Applications include medical illustration, pathology research, surgical simulation, and general scientific education.

Color photographic volume data greatly simplifies part of the difficulty in creating realistic volume rendered images: determining the appropriate color value for each voxel in the data set. Photographic volume data also offers a challenge to traditional volume rendering techniques: determining the opacity for each voxel of the data set. In

traditional direct volume rendering, an image is produced from a volume of scalar data, using transfer functions from scalar value to color and opacity. The design of effective color and opacity transfer functions from scalar values has been the subject of substantial research over the past 10 years (e.g., [8], [9], [16]), with the design of the color transfer function (1D to 3D mapping) often being much more difficult than the design of an effective opacity transfer function (1D to 1D mapping). In contrast, volume rendering from photographic volume data sets reverses the difficulty of transfer function design. Photographic volume data sets need an opacity transfer function from vector color data to scalar opacity data (3D to 1D), a process complicated by the nonlinear nature of color spaces. By contrast, for photographic data, the design of an appropriate color transfer function is not generally required since the color of each voxel is already known.

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

Volume rendering research often applies differential geometry of scalar data fields to create normals and information about probable boundaries. This information is then used to shade, light, and composite shapes, objects,

- D.S. Ebert is with the School of Electrical and Computer Engineering, Purdue University, 1285 EE Building, West Lafayette, IN 47907. E-mail: ebertd@purdue.edu.
- C.J. Morris is with the Visual Technology Group, IBM T.J. Watson Research Center, 19 Skyline Dr., Hawthorne, NY 10532. E-mail: cjmorris@us.ibm.com.
- P. Rheingans is with the Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. E-mail: rheingan@cs.umbc.edu.
- T.S. Yoo is with the Office of High Performance, Computer and Communication, National Library of Medicine, National Institutes of Health, Building 38A, Room B1N30, 8600 Rockville Pike, Bethesda, MD 20894. E-mail: yoo@nlm.nih.gov.

Manuscript received 2 Oct. 2000; revised 16 July 2001; accepted 24 July 2001. For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number 112927.

and features embedded within the data. Volume rendering vector data is not well understood: Calculating boundary information and normals to probability surfaces from vector fields is an area of open research. Color fields add the complication that the spaces in which they live are nonlinear with respect to human perception. Thus, direct rendering of photographic volumes requires: 1) reparameterization of the native color values to a more perceptually linear space where differential geometry may be usefully applied and 2) the exploration of effective mappings of 3-valued color vectors to a scalar opacity field.

We present our approach for volume rendering directly from photographic data without the need for a secondary scalar volume to indicate density and show examples of applying these techniques to direct volume rendering of the Visible Human Male data set. Direct volume rendering from the photographic volume eliminates the need to register volumes from different sources and enables the display of internal volume structure along with material boundary information. We describe our experiments using different color spaces and color components for determining the volume opacity and evaluate the potential of these techniques for depicting anatomical structures.

2 PHOTOGRAPHIC VOLUME DATA

Photographic color data is becoming increasingly important as a volume information representation. While the acquisition of such data often requires the loss of physical integrity of the sample, tomographic sectioning has been important in anatomy and in pathology. With the growing capability of aggregating multiple photographic cutplane images into 3D volumes, these techniques are growing in importance in other fields. Limitations of MR imaging in generating adequate resolution led the Whole Frog Project at Lawrence Berkeley Laboratory to create an entire frog data set using cryosection [22]. The use of this data set is growing as a basis for teaching dissection in biology. The Laboratory for Neurological Imaging at UCLA commonly uses cryosectioning to gain the resolution and contrast required for their intricate analysis of the brain's pathways [29]. Commercial groups now offer mechanical sectioning technology that is used in volume data analysis for geology, medicine, and semiconductor manufacturing [12]. The National Library of Medicine produced one of the most important examples of data acquired through photographic tomographic imaging: the Visible Human data sets.

The Visible Human Project was formed to explore the use of digital imaging technology in modern anatomy research and education. Data from two subjects, one male and one female, were collected through a variety of methods, including the conventional radiological techniques of X-ray CT studies, magnetic resonance imaging (MRI), and plain film radiographs. In addition to these conventional clinical studies, the subjects were frozen and sectioned at 1 mm (male subject) and 1/3 mm (female subject) intervals. The exposed surfaces were photographed with 35 mm and 70 mm film and digitized with an electronic camera. Image acquisition was carefully performed and the resulting data is one of the most complete anatomical studies ever performed [27]. Each slice of the

digital cryosection data was acquired with a raster resolution of $2,048 \times 1,216$ pixels with a horizontal field of view of 25 inches. Voxel dimensions are $0.32 \times 0.32 \times 1$ mm in the male data set. The resulting data sets are 14 gigabytes and 40 gigabytes, respectively.

3 RELATED WORK

Visualization of volume data remains a challenge across multiple scientific disciplines. Several approaches to illuminating the internal structure of volume information have been explored as part of the growing effort in visualizing complex, high-dimensional data. From its inception, volume visualization employs a wide array of methods and techniques. In our work, we assign optical properties to volume information based on changes and gradations of the color data. Alternate approaches employ similar mathematics and are based on the same fundamental observations.

3.1 Volume Visualization

Direct volume rendering as an image generation and reconstruction technique has a rich history in computer graphics [18], [3], [14]. Early papers included medical data acquired using X-ray CT scans, assigning opacity on the basis of X-ray attenuation. Initial implementations and much of the derivative work used pseudocolor to improve the visualizations produced through volume rendering. The recent availability of hardware 3D texture capabilities gives rise to a volume rendering mechanism which more easily accommodates color data [1]. However, these methods must still generate an opacity transfer function to enable correct compositing of the image planes. One technique for doing so involves using scalar data at a particular voxel and its surrounding neighbors to produce a local gradient vector for each voxel [18]. Subsequently, this local gradient vector and its magnitude are then used to calculate an opacity value for the particular voxel.

Researchers working with the Visible Human data have analyzed the color gamut of the male data set in their work on photorealistic volume rendering and virtual dissections [15], [21]. Since reflectance of the light rays is often based on gradients measured in the object volume, careful consideration of the color spaces involved should be part of the work. Sapiro and Ringach showed that selection of the color space can make dramatic differences when attempting nonlinear image processing [24].

Other techniques for extracting visual information from volumes include isosurface extraction techniques such as the Marching Cubes algorithm [20]. Early work by Lorensen on the Visible Human data either used the X-ray CT data or separated the red channel of the RGB images to generate isosurfaces from scalar rather than multivalued data [19]. Similar work by the Vesalius group at Columbia University extracts isosurfaces from the color data and later uses the original color volume as a solid texture to apply color information to the extracted polygonal surfaces [13]. Both of these approaches have involved strictly surface, rather than volume, rendering.

3.2 Transfer Functions versus Segmentation

Visualization of color data or other multivalued data can be accomplished as part of a concatenation of filtering, segmentation, and rendering steps. Working from segmented data simplifies the problems of opacity assignment, but requires an investment in segmentation, often a labor-intensive process of hand segmentation. Renderings of this type based on hand-segmented volume color data of the VHP thorax were published by Zhou and Henderson [32]. Other work in progress by Takashi et al. describes a pipeline of filtering and morphological processing as part of a segmentation step that is followed by visualization [28]. Color segmentation has been explored independently of visualization as part of computer vision research. Sapiro applied the finite element techniques of active contours pioneered by Witkin, Kass, and Terzopolous to color images by adapting color metrics to the systems of differential equations that underly their approach [25]. He created a more comprehensive treatment of multivalued image analysis as part of a broader text [26].

Color data has a particular structure and is a subset of a wider domain of multivalued data. A variety of more generic multivalued nonlinear filtering techniques have been the subject of many research projects. Nonlinear filtering of MRI data has been developed by Gerig et al. for surgical planning [10]. Also related, Laidlaw et al. have explored the classification of multivalued MRI data using within-pixel statistics to adjust for partial voluming problems as part of a broader effort to create geometric models from volume medical scans [17]. Color and color spaces, however, provide a specific perceptual model for the detection of objects, boundaries, and features. Our research concentrates on utilizing these characteristics as part of the visualization effort.

We reported early results in volume rendering the brain of the Visible Human Male in [4]. Some of our earlier work on volume rendering of photographic data and the surface-based techniques of Lorensen and of the Vesalius group share the drawback that they both require a segmentation of the data before rendering. Because of the difficulties in segmenting color data, visualization projects often resort to hand segmentation, a labor intensive interactive process requiring input from experts in anatomy and physiology. This technique generates a handcrafted volume of interest or a binary mask to delineate the regions to be rendered. However, these partitions are seldom created with rendering issues in mind, leading to sampling artifacts, partial voluming errors, and other difficulties that must be overcome later in the visualization process. Moreover, committing to a particular segmentation inhibits the exploration of a volume data set. The task becomes the visualization of the *segmentation* rather than of the information within the data. Our current work focuses on techniques necessary for direct rendering of color data.

By contrast, direct visualization of volume data can be accomplished without first performing a segmentation. Transfer functions are used to map image properties to illumination and rendering characteristics, including opacity, color, and texture. A survey of transfer function techniques can be found in Pfister et al.'s report on the

Visualization panel [23]. While many of these methods share similar mathematics with nonlinear color filtering and color segmentation, the goal of applying transfer functions is to visualize the data without creating a fixed model of the structure of the data. As mentioned previously, our work in deriving density/opacity values from photographic data is different from traditional transfer function approaches in the following ways:

- Our source data is three-component data.
- The source data domain is perceptually nonlinear.
- Color data has unique, inherent structural characteristics that we can utilize.

Therefore, we have chosen to concentrate on finding a series of transfer functions that is well-suited to the problem domain and not on generating general purpose vector or scalar transfer functions.

4 COLOR IN PHOTOGRAPHIC VOLUMES

Although color information is available in the photographic description of each voxel, the data does not contain any opacity information or information on the viewpoint-dependent reflective or light transmission properties of the voxels. The color values of each voxel describe the total reflection from that voxel for the specific set of lighting and viewing parameters where the voxel is lit and viewed from directly above. We have chosen to make the simplifying assumption that this total reflection is just diffuse reflection.¹ Therefore, we can use the same color as the reflection from other viewpoints.

4.1 Color Spaces for Volume Rendering

The original digital photographic data is 8-bit R, G, and B color values (see Fig. 1 for the color slice and each individual component). While structures of interest are readily apparent to the anatomically trained eye, automatically identifying detail to be emphasized in a volume rendering is more difficult. We consider the goal of the opacity transfer function to be to capture as much of the anatomical structure visible to the human eye as possible. We base our approach on the characteristics of human vision, not because that is the only possible approach, but because human vision demonstrably works on these images. Our approach draws from the functions of very low-level vision, concentrating on basic changes in the visual field. Other approaches including elements from higher-level vision would be interesting future extensions, but would necessarily start from the basic differences we attempt to capture here.

Since the RGB color space does not correspond to human visual perception, we chose to explore the CIE $L^*u^*v^*$, abbreviated CIELUV, color space to obtain a perceptually uniform representation of the color volume. A perceptually uniform color space has the characteristic that equal distances in the color space correspond to equal perceptual differences, at least for reasonably small distances. This

1. While using a correct bidirectional reflection distribution function for the voxel would yield more accurate results, determining the BRDF from photographic data is still an open research question. Considering the reflection to be simply diffuse results in a reasonable first approximation.




Fig. 1. (a) Reference rendering of single slice. (b) Intensity is R color component. (c) Intensity is G color component. (d) Intensity is B color component.

allows us to computationally estimate the perceptual difference between two image samples by simply computing the Euclidean difference in color space. Using a perceptually uniform color space for opacity transfer function calculations allows us to emphasize those features which are noticeable in the photographs, creating a volume rendering more consistent with features detected by human vision in the photographs than that using a device-derived color space, such as RGB, which would overemphasize changes in blue content and bright regions. The CIE $L^*u^*v^*$ color space, in particular, also offers the advantage that chromatic and achromatic components of color are described by orthogonal color space dimensions. This feature allows us to experiment with biased weighting of the chromatic and achromatic color components, just as the human visual system performs certain scene understanding tasks with segregated achromatic and chromatic color information.

The photographic image of a slice from the Visible Male data set and each CIE $L^*u^*v^*$ color component can be seen in Fig. 2. The L^* component corresponds to the perceptually linear lightness of the color, so it has large values in light areas of the volume. For human tissue, this is largest for bones, skin, fat, and light colored organs (e.g., brain). The u^* component captures chromatic changes in roughly red-green colors. Therefore, this component is good in capturing changes in the "redness" of tissues (e.g., muscle to bone). The v^* component captures color changes in the yellow-blue dimension. For most human tissue, this would not seem to be very useful in distinguishing tissue types. In these images, it is evident that there is much more variation in the u^* (red-green) color component than in the other two components. The gray level changes in the images create a three-dimensional effect in the slice since our eyes tend to interpret shading variation as the result of shape variation.

This effect is strongest in the intensity image for the u^* component since the red-green color changes in the photograph correspond strongly to the boundaries of the muscles and of the fat.

Fig. 3a illustrates the inadequacies of the RGB color space for modeling perceptually meaningful color metrics. This figure shows the cumulative perceptual difference between adjacent colors in ramps of the individual RGB color components. Notice that none of the ramps approach linearity, understating changes at low intensities and overstating those at high intensities. Notice also how increments in the blue content cause relatively small perceptual differences.

Other common color models that might be used include those organized to match artistic notions of color based on hue and lightness, such as HLS (Hue-Lightness-Saturation) and HSV (Hue-Saturation-Value). Like CIE $L^*u^*v^*$, these spaces attempt to separate the chromatic and achromatic components of color, potentially enabling a task-based decomposition of color. Since linear transformations of RGB, HLS, and HSV are also device-dependent color models, they display substantial variation in brightness displayed for constant values of L or V. A clear example of this can be seen by viewing a constant L slice of a typical color picker under black and white viewing conditions, invariably showing brightness to be far from constant. The primary advantage of these color models is their enabling of more intuitive color naming for humans, rather than more accurate color difference measures. Fig. 3b shows the cumulative perceptual difference between the individual HSV color components. For the hue curve, full saturation and value have been used. The saturation and value curves use a red hue and a maximum value of the other unvarying component. The value curve shows that differences at low intensities are understated, while those at high intensities




Fig. 2. (a) Reference rendering of single slice. (b) Intensity is CIE L color component. (c) Intensity is CIE u color component. (d) Intensity is CIE v color component.

are overstated. The wiggles in the hue curve show that changes in hue produce different amounts of perceptual differences for different hue ranges. The saturation curve shows how little perceptual difference small changes in saturation have. These curves illustrate that HSV is not much better suited than RGB to capture the perceptual differences present in color photographic data. Alternatively, the use of a more complex color appearance model which considers effects of adaption and context [7], rather than simply colors in isolation against a neutral background, would be interesting, but is beyond the scope of this research.

Color space conversions were performed using the methods of Hall [11]. Because precise specifications of the color primaries of the image data were not directly available, we approximated them by the NTSC standard primaries. Specification of the image data color primaries is indirectly available from the color calibration card included in each photograph. Analysis of the RGB values for these physically measurable colors could provide the required calibration information. Although deviations of actual primaries from these are expected to be modest, this approximation does not guarantee the correct absolute CIE $L^*u^*v^*$ coordinates for voxels, compromising the device independence of the color space. For color reproduction applications, this would be a problem, but, for the purposes of detecting color structure in an image, it is not. Fortunately, the approximation should preserve the relationships between points in the color space, which is our primary concern since relative judgment, rather than absolute, is the basis for almost all perceptual processes.

5 COLOR DISTANCE GRADIENTS

Volumetric data is often a series of discrete samples acquired on uniform intervals from physical objects. The Visible Human Project data are rectilinear arrays of color values sampled directly from a human subject. Since changes between neighboring voxels are described mathematically as a color derivative, we use terminology and techniques from differential geometry to aid in manipulating the volume data.

We follow a model similar to the color difference models described and developed independently by Sapiro [25] and Cumani [2]. Since our methods use the gradient and its magnitude as the input for an optical transfer function rather than a binary or probabilistic segmentation problem, we are able to defer some of the more difficult issues of classification to the human observer. Our treatment, therefore, begins with the same fundamental analysis of color spaces, deriving a general form for color gradient and color gradient magnitude.

For color photographic volume data represented in CIELUV space, every voxel location is mapped to a color, represented by three separate values. Given $color_vol[x][y][z]$, a volume of color datapoints, imagine its image generating function $C(x, y, z)$ such that, for every voxel location $[x_i][y_i][z_i]$,

$$color_vol[x_i][y_i][z_i] = C(x_i, y_i, z_i) = (L_i, u_i, v_i). \quad (1)$$

As we mentioned before, CIELUV space has reparameterized color space so that, for unsaturated colors, the perceptual distances between two colors is, to a good approximation, proportional to the length of the straight line joining them. In other words, the color space is approximately locally linear. We, therefore, consider color values in CIELUV space as *color vectors*. Given this




Fig. 3. Cumulative difference curves of color components of (a) RGB color model, (b) HSV color model.

assumption, color differences can be measured with respect to spatial location. For instance, consider the difference between a color voxel value $C(x_i, y_i, z_i)$ and the color from a neighboring location measured at a distance, h , along the x direction, $C(x_j, y_j, z_j)$, where $x_j = x_i + h$, $y_j = y_i$, and $z_j = z_i$.

The color difference vector is expressed as:

$$\begin{aligned} C(x_i, y_i, z_i) - C(x_j, y_j, z_j) \\ = C(x_i, y_i, z_i) - C(x_i + h, y_i, z_i) \\ = (L_i, u_i, v_i) - (L_j, u_j, v_j). \end{aligned} \quad (2)$$

Note that the difference of two adjacent colors is itself a color vector. That is:

$$(L_i, u_i, v_i) - (L_j, u_j, v_j) = (L_i - L_j, u_i - u_j, v_i - v_j). \quad (3)$$

As h is reduced and made infinitesimally small, we derive the partial derivative of $C(x, y, z)$ with respect to x . That is,

$$\begin{aligned} \frac{\partial}{\partial x} C(x, y, z) \\ = \lim_{h \rightarrow 0} C(x, y, z) - C(x + h, y, z) \\ = \left(\frac{\partial}{\partial x} L(x, y, z), \frac{\partial}{\partial x} u(x, y, z), \frac{\partial}{\partial x} v(x, y, z) \right). \end{aligned} \quad (4)$$

The combined partial derivatives of $C(x, y, z)$ in the x , y , and z directions results in $\nabla C(x, y, z)$, the gradient of $C(x, y, z)$. Since each partial derivative is a color vector, $\nabla C(x, y, z)$ is a tensor, a vector of three color vectors.

$$\nabla C(x, y, z) = \left(\frac{\partial}{\partial x} C(x, y, z), \frac{\partial}{\partial y} C(x, y, z), \frac{\partial}{\partial z} C(x, y, z) \right). \quad (5)$$

We can evaluate the magnitude of $\nabla C(x, y, z)$ by taking its tensor dot product. This is accomplished by multiplying $\nabla C(x, y, z)$ with its transpose $\nabla C(x, y, z)^T$, resulting in Λ , the "magnitude matrix." Mathematically:

$$\Lambda = (\nabla C(x, y, z))(\nabla C(x, y, z)^T) = \nabla C(x, y, z) \cdot \nabla C(x, y, z). \quad (6)$$

We can extract useful geometry from the tensor gradient magnitude matrix Λ . At this point, we diverge from the more generic problem of object recognition, segmentation, and filtering to concentrate on the generation of optical transfer functions for volume rendering. Filtering and curve and surface evolution methods for segmentation are founded on nonlinear partial differential equations. The stability of these differential methods depends on the regularity of the discrete approximations created by the voxel grid to a continuous, differentiable image representation. For our purposes, however, the transfer functions simply treat large color gradients uniformly as discontinuous areas and assign an optical property such as high opacity to that location. Given a locally linear color space such as $L^*u^*v^*$, Euclidean distances can be used as a relatively good approximation to the color gradient magnitude in smoothly varying areas of the volume. At discontinuities, $L^*u^*v^*$ color differences between voxels can no longer be approximated by a linear function; however, the nonlinear behavior of the transfer function mitigates the need for accurate measurement of the gradient magnitude at boundaries, relying instead on the detection of the existence, location, and direction of the boundary. Geometric direction, not color orientation, is required for illumination and shading. Given these parameters (approximate local linearity of color differences in smooth regions and requiring only direction and location but not accurate color gradient magnitude measurement across boundaries), we concentrate on uses for the diagonal elements of Λ to capture these properties. In particular, we examine their square roots:

$$grad.x = \sqrt{\Lambda_{11}} = \sqrt{\frac{\partial}{\partial x} C(x, y, z) \cdot \frac{\partial}{\partial x} C(x, y, z)}, \quad (7)$$

$$grad.y = \sqrt{\Lambda_{22}} = \sqrt{\frac{\partial}{\partial y} C(x, y, z) \cdot \frac{\partial}{\partial y} C(x, y, z)}, \quad (8)$$

$$grad.z = \sqrt{\Lambda_{33}} = \sqrt{\frac{\partial}{\partial z} C(x, y, z) \cdot \frac{\partial}{\partial z} C(x, y, z)}. \quad (9)$$

These values reflect the color changes along the cardinal Cartesian directions and their influence on the color distance gradient magnitude.² We call the vector quantity $(grad.x, grad.y, grad.z)$ the *color distance gradient*. This vector value is not a *true* gradient since it does not capture absolute direction (there are sign ambiguities in each of the cardinal directions) nor does it capture the absolute magnitude of the matrix Λ . However, as an abstraction of the magnitude matrix, it can be gainfully used in the generation of opacity transfer functions.

To measure these geometric values on discrete voxel data, we use the common method of central differences to compute approximate derivative values. These calculations were performed both on the normalized CIELUV color data as well as on the original RGB values and the results were compared. We compute the discrete approximation as:

2. The variation of individual color components along each Cartesian direction is captured in the off-diagonal elements of the matrix.

$$\begin{aligned} grad.x &= color_distance(color_vol[x-1][y][z], \\ &\quad color_vol[x+1][y][z]) \\ grad.y &= color_distance(color_vol[x][y-1][z], \\ &\quad color_vol[x][y+1][z]) \\ grad.z &= color_distance(color_vol[x][y][z-1], \\ &\quad color_vol[x][y][z+1]), \end{aligned} \quad (10)$$

where, for two voxels $v1$ and $v2$,

$$\begin{aligned} LUV_distance(v1, v2) \\ = \sqrt{(v1.l - v2.l)^2 + (v1.u - v2.u)^2 + (v1.v - v2.v)^2} \end{aligned} \quad (11)$$

$$\begin{aligned} RGB_distance(v1, v2) \\ = \sqrt{(v1.r - v2.r)^2 + (v1.g - v2.g)^2 + (v1.b - v2.b)^2}. \end{aligned} \quad (12)$$

In both cases, we use the Euclidean color space distance between colors as a scalar metric describing the color difference. For CIELUV, this is equivalent to the CIE 1976 ($L^*u^*v^*$)-space color-difference formula [30] and is perceptually sound since Euclidean distances correspond to perceptual differences. For RGB colors, this is not guaranteed to give an accurate measure of perceptual difference, due to the nonlinearities of the color space, but no better metric is available. The color difference metric used is unsigned distance, giving the sizes but not direction of the difference. Because the difference metric is scalar, it captures the perceptibility of change in that region, but not the nature. Specifically, the metric is not sensitive to differences in particular color space components, only the total difference.

Fig. 4 shows the color difference gradient vectors in the $L^*u^*v^*$ color model displayed using the Stereoscopic Field Analyzer, a glyph-based volume visualization environment [6]. A magnified view of a portion of the slice (muscle tissue in the lower lefthand corner of the slice) is represented by both RGB gradient vectors and $L^*u^*v^*$ gradient vectors (Fig. 5a and Fig. 5b, respectively). From these two images, differences in the two gradient spaces could be seen rather clearly. The $L^*u^*v^*$ gradient vectors are not only greater in size (corresponding to the magnitudes), but, in areas where the RGB gradient vectors appear consistent and uniform, the corresponding $L^*u^*v^*$ vectors vary in orientation, capturing more detail.

Alternatively, the gradient of a single component of the color could also be used to emphasize information well-captured by that color component, such as CIE u^* red-green color variation orientation. This same result can be achieved by applying gradient-based shading techniques [18] to the color component scalar fields described in the previous section.

6 OPACITY TRANSFER FUNCTIONS

A major challenge with volumetric photographic data is determining an appropriate transfer function from the vector valued voxel color to the voxel density/opacity. We have explored several different opacity transfer functions using both RGB and CIE $L^*u^*v^*$ color spaces.



Fig. 4. Visible Human Slice rendered in SFA with vector glyphs. The X, Y, and Z components of each glyph correspond to the X, Y, and Z components, respectively, of the $L^*u^*v^*$ gradient at that point.

For each candidate transfer function, the density of each voxel for volume rendering was computed using the selected transfer functions. The volume was then rendered using a modified volume ray tracer that uses atmospheric accumulation, attenuation, illumination, and shadowing [31], [5]. The following simple, final opacity transfer function was used in the rendering process to allow control over the opaqueness of the volume:

$$renderedopacity = (voxel_opacity * scalar)^{exponent}. \quad (13)$$

The *scalar* variable allows the overall opaqueness of the volume to be controlled, while the *exponent* variable allows the sharpness of the opacity fall-off to be easily controlled (larger *exponent* produces sharper opacity transitions).

An effective opacity/density transfer function for anatomical photographic data should capture features and structures within the human body based on the color photograph data. We performed an initial analysis using




Fig. 5. Close up of a muscular area in the lower lefthand section of the slice in Fig. 4 with the X, Y, and Z directions of the vector glyphs corresponding to (a) the X, Y, Z components, respectively, of the RGB gradient at that point and (b) the X, Y, and Z components, respectively, of the $L^*u^*v^*$ gradient at that point.




Fig. 6. Slice rendered in SFA with the opacity set to (a) the RGB gradient magnitude, (b) the $L^*u^*v^*$ gradient magnitude.




Fig. 7. Composited renderings where opaque (yellow/green) top layer is computed from opacity equal to (a) RGC gradient magnitude, (b) CIE $L^*u^*v^*$ gradient magnitude, (c) RGB gradient dot product, (d) CIE $L^*u^*v^*$ gradient dot product.

one slice from the thorax section of the Visible Male data set (slice vm1300) to explore the effectiveness of several transfer functions. This slice was chosen because it contained a significant amount of muscle, bone, fat, and other tissues. A constant opacity rendering of this slice can be seen in Fig. 2a. We have explored eight different color to opacity transfer functions.

6.1 Color Components

To see the chromaticity and luminance differences within the data, we developed transfer functions that set the density of the data set equal to the separate CIE L^* , u^* , and v^* color components. The resulting images can be seen in Fig. 2b, Fig. 2c, and Fig. 2d, respectively. As noted earlier, the u^* component is much more successful at capturing tissue boundaries and internal structures.

6.2 Color Distance Gradient Magnitudes

In volume rendering, it is often useful to accentuate the transitions or boundaries between regions that have contiguous colors, gray-levels, or other properties. Unlike the scalar gradient magnitude values that arise from gray-level volume data, the tensor gradients of color photographic data are not easily used to generate opacity values. One approach is to take the magnitude of the color distance gradient vector or $\|(\text{grad}.x, \text{grad}.y, \text{grad}.z)\|$.

This measurement will make the transitions between regions more opaque when adjacent voxel colors are more distant. We apply these measurements to both the RGB color volume data arrays and the normalized CIELUV volume data arrays, generating the color distance gradient $rgb_grad(\text{voxel})$ and $luv_grad(\text{voxel})$, respectively. The resulting opacity transfer functions are the following:³

$$\text{voxel_opacity} = \|RGBgrad(\text{voxel})\| \quad (14)$$

and

$$\text{voxel_opacity} = \|LUVgrad(\text{voxel})\|. \quad (15)$$

Fig. 6 shows the resulting SFA images achieved when the volume density is set to RGB and CIE $L^*u^*v^*$ gradient magnitudes. Fig. 7a and Fig. 7b, respectively, are the corresponding volume rendered images. In these images, the gradient magnitude image is composited on top of the original photographic slice with high opacity for larger gradients to illustrate the features of the data that each technique was able to capture. The CIE $L^*u^*v^*$ color distance gradient captures more of the visible detail within

3. These opacity transfer functions correspond to the square root of the trace of the color distance gradient magnitude matrix, Λ . For diagonalized matrices, the trace is actually the gradient magnitude.




Fig. 8. Slice rendered in SFA with the opacity set to (a) the RGB color distance gradient dot product, (b) the $L^*u^*v^*$ color distance gradient dot product.

the data in certain areas, such as near the humerus in the right of the image. When compared with the RGB color distance gradient, this higher level of detail is due to the closer correspondence to human visual characteristics of the $L^*u^*v^*$ colorspace. Although these particular transfer functions are somewhat effective and are able to pull out most of the features (skeletal structure, fatty tissue, etc.), there is still a significant amount of detail that was not extracted, prompting further investigation.

6.3 Color Distance Gradient Dot Products

The color distance gradient magnitude only captures changes in the length of the color distance gradient vector, not changes in its spatial orientation. Since a gradient vector can be computed for each voxel, a normal and its corresponding tangent plane can be associated with each voxel, denoting the orientation of the tangent space that passes through that location. In areas of smoothly changing orientations, edge information is dominated by the color distance gradient magnitudes. However, in turbulent or highly textured regions, the orientations of nearby normal vectors become highly varying. To capture information regarding the relative smoothness or roughness of volumetric textural features, we can compute the color distance gradient at a voxel and its angular difference with its neighboring adjacent voxel gradients (in our discrete implementation, we only compare the color distance gradient at a voxel with its six neighbors in the x , y , and z directions). The angular separation between the color difference gradients of two neighboring voxels is inversely proportional to the dot product of their normalized color difference gradient vectors. Therefore, to highlight areas of large angular variation of color difference gradient vectors, the following transfer function can be used:

$$voxel_opacity = \max_{i=1}^6 (1 - Grad(voxel) \cdot Grad(neighbor_i)), \quad (16)$$

where

$$Grad(voxel) = \frac{grad(voxel)}{\|grad(voxel)\|}$$

and $\|grad(voxel)\|$ can be either the magnitude of the RGB color distance gradient or the magnitude of the CIE $L^*u^*v^*$ color distance gradient.

This function is good for highlighting the largest angular variation, regardless of the size of the color distance gradient. A more common approach would be to also weight the opacity by the length of voxel color distance gradient, yielding the following transfer function:

$$voxel_opacity = \max_{i=1}^6 [(1 - Grad(voxel) \cdot Grad(neighbor_i)) * \|grad(voxel)\|]. \quad (17)$$

This transfer function will highlight areas with the large color distance gradients that have the largest angular variation with their neighboring voxels. This function will, therefore, be useful for highlighting the largest color changes.

To highlight small scale changes within the color data (e.g., muscle fibers) as well as larger scale changes, the minimum angular color distance gradient change times the gradient magnitude is very useful. This formula enhances areas where there are large gradients but small angular color changes with one of the six adjacent voxels. Therefore, it will highlight oriented patterns within the color data. The resulting transfer function is the following:

$$voxel_opacity = \max_{i=1}^6 [(Grad(voxel) \cdot Grad(neighbor_i)) * \|grad(voxel)\|]. \quad (18)$$

The results of using the dot product of a voxel's gradient vector and that of the neighbor most similarly oriented (coupled with the gradient magnitude), in RGB and $L^*u^*v^*$ space, can be seen in Fig. 7c and Fig. 7d, respectively. The corresponding images are displayed in Fig. 8. These transfer functions capture much more detail than the gradient magnitude alone. This improvement in detail is due to the fact that areas where the gradient magnitudes are large (i.e., skeletal structure and fatty tissue, as seen previously), as well as areas where the local gradient vectors are similarly oriented (i.e., along muscle fibers, tissue edges), are rendered opaque. Using a transfer function which emphasizes places where both indicators are large, helps to pick




Fig. 9. Volume rendering of abdominal region of Visible Man data set (slices vm1600 to vm1662). (a) Opacity set to CIE $L^*u^*v^*L^*$ color component. (b) Opacity set to CIE $L^*u^*v^*u^*$ color component. (c) Opacity set to CIE $L^*u^*v^*v^*$ color component. (d) Opacity set to RGB colorspace distance gradient magnitude. (e) Opacity set to CIE $L^*u^*v^*$ colorspace distance gradient magnitude.

out fine continuing structures. Edges between tissue and bone were brought out by the previous methods; however, this method was much more effective in pulling out individual muscle fibers and other less obvious edges, creating a more detailed and informative rendering. The CIE $L^*u^*v^*$ gradient dot product transfer function also captures more visible detail within the muscles than the corresponding RGB gradient calculation. As before, the

level of detail achieved by using the $L^*u^*v^*$ colorspace, as compared to the RGB colorspace, can be attributed to the higher level of variance within the $L^*u^*v^*$ colorspace.

All of the above color distance gradient dot product transfer functions could be applied directly to the actual RGB and CIE $L^*u^*v^*$ color components, rather than the color difference gradient vectors. However, the resulting transfer functions are not as effective in capturing tissue




Fig. 10. Volume rendering of 64 slices. (a) Opacity determined by CIE $L^*u^*v^*u^*$ color component. (b) More opaque rendered, cut away view with opacity determined by CIE $L^*u^*v^*u^*$ color component. (c) Cut away view with opacity from rgb color distance gradient and product. (d) Cut away view with opacity from CIE $L^*u^*v^*$ color distance gradient dot product. (e) More transparent cut away view with opacity from RGB color gradient dot product. (f) More transparent cut away view with opacity from CIE $L^*u^*v^*$ color gradient dot product.




Fig. 11. volume rendering of photographic Visible Human Male head, without and with enhancement of low gradient areas from side and top views. Transfer function uses the CIE $L^*u^*v^*$ color distance gradient to determine the opacity.

boundaries and internal tissue structures. If normalized color space color vectors are used, differences within each color component are lost (e.g., bright red and dark red). If unnormalized colorspace color vectors are used, colors with the largest component colors will be more opaque and differences in the small color component colors will be lost.

7 COMPARISON OF TRANSFER FUNCTIONS ON VOLUME OF DATA

Volume renderings of an abdominal section of the Visible Male data set (slices vm1600-vm1663) using each CIE $L^*u^*v^*$ component, the RGB color distance gradient magnitude, and the CIE $L^*u^*v^*$ color distance gradient magnitude can be seen in Fig. 9. In this particular region, there are many organs of interest (i.e., kidneys, intestines, stomach, etc.) and the transfer functions were able to generate renderings that maintained a clear delineation between these varied organs, demonstrating their usefulness. From comparing the images in this figure, it is clear that the CIE $L^*u^*v^*$ colorspace gradient is much more

effective than RGB for capturing more detail within the muscles, kidneys, etc.

We also applied these techniques to a thorax and head section of the Visible Male data set. Fig. 10 shows the results of the direct volume rendering using the CIE $L^*u^*v^*$ u^* color component and the gradient dot product transfer functions applied to 64 slices of the Visible Human Male photographic database (slices vm1300-vm1364). Fig. 10a shows a rendering of data as seen from the outside, which uses a more transparent scaling in the rendering process. Fig. 10b uses the same transfer function, but contains a cut-away view to display the interior and uses a more opaque scaling. Both of these images highlights the muscles (e.g., deltoids) and their internal muscle fibers clearly. The spinal column can also be seen in the back of the images.

Fig. 10c, left image, and Fig. 10d, right image, show the results of using the RGB and CIE $L^*u^*v^*$ color space gradient dot product, respectively, for calculating the voxel density. Both of these images display the bones, muscles, and fat tissue effectively. The CIE $L^*u^*v^*$ color space dot product captures more muscle detail, giving Fig. 10d more




Fig. 12. Sagittal view of volume rendering of a $1,024 \times 512$ centered subset of the upper abdominal region of the Visible Man data set (slices vm1300-1363). The opacity is calculated from the CIE $L^*u^*v^*$ color distance gradient showing the spine, diaphragm, and interior structure of the heart, especially the left ventricle.



Fig. 13. Sagittal view of volume rendering of a $1,024 \times 512$ centered subset of the upper abdominal region of the Visible Man data set (slices vm1300-1363). The opacity is calculated from the CIE $L^*u^*v^*$ color distance gradient minimum dot product, showing the spine, diaphragm, and interior structure of the heart. The rendering parameters were chosen to highlight liver and heart muscle interiors more than the previous image (exponent = 0.7).

red tones. The difference between the color spaces for capturing muscle detail can be seen more clearly in Fig. 10e and Fig. 10f, where a more transparent rendering was produced. Again, the CIE $L^*u^*v^*$ color space image (Fig. 10f) shows more muscle fiber, especially in the deltoid region.

A volume rendering from the photographic data of the Visible Human Male head using the CIE $L^*u^*v^*$ color distance gradient magnitude is presented in Fig. 11. Fig. 11a contains a side view image generated by a higher exponent value (exponent = 1.2), making the interior tissue structures more transparent. Fig. 11b shows the same view image generated with a lower exponent (exponent = 0.9) to increase the importance of gradients within tissue types. Both of these images show tissue boundaries between gray matter and the Corpus Collosum, a lateral ventricle, portions of the skull, and the sinus cavity. Fig. 11c and Fig. 11d are top view images demonstrating the range of images that can be generated by varying the exponent to show all internal tissues (Fig. 11c, exponent = 0.5) or to show primarily tissue boundaries (Fig. 11d, exponent = 1.0).

Fig. 12 shows a sagittal view volume rendering of a $1,024 \times 512$ centered subset of the upper abdominal region of the Visible Human Male data set (slices vm1300-1363). The opacity is calculated from the CIE $L^*u^*v^*$ color distance gradient minimum dot product showing the spine, aorta, diaphragm, and interior structure of the heart, especially the left ventricle. Rendering parameters were chosen to sharpen tissue boundaries (exponent = 1.3). Fig. 13 uses the same opacity transfer function to highlight the spine, diaphragm, and interior structure of the heart. The rendering parameters were chosen to highlight liver and heart muscle interiors more than the previous image (exponent = 0.7).

8 SUMMARY AND FUTURE DIRECTIONS

We have proposed a basic framework for addressing the rendering challenges presented by photographic volume data. The images presented in this paper show that effective volume rendering can be performed from photographic volume data sets without the need of auxiliary data sets for density information. The CIE

$L^*u^*v^*$ colorspace provides useful information for capturing anatomical structures within the Visible Human photographic data set due to its correspondence to human visual system characteristics. The CIE $L^*u^*v^*$ u^* component is very effective in rendering muscles, while the CIE $L^*u^*v^*$ color distance gradient dot product transfer function is capable of capturing both muscle detail as well as bone, fat, and other tissue detail. This is a result of heavily weighting areas of large magnitude, as well as areas of consistent gradient orientation.


As effective as these methods are, many open issues remain. These issues include color calibration and internal registration methods, the optimal weighting of colorspace dimensions in gradient calculations, more realistic reflectance models, and determination of segment-based densities directly from the photographic volume. Furthermore, it should be mentioned that the colorspace components that proved to be optimal for volume rendering of the Visible Human Male may be specific to human data. The weightings and calibrations of subsequent opacity transfer functions may have to be adjusted for various data sets, for example, data that is nonorganic in nature. We feel that our general approach will still be able to guide us along the right path in determining these changes. As a result, we plan to explore the effectiveness of these techniques for capturing features in other photographic volume data sets.

ACKNOWLEDGMENTS

This work supported in part by the US National Science Foundation under Grant Nos. ACIR 9996043 and ACIR 9978032.

REFERENCES

- [1] B. Cabral, N. Cam, and J. Foran, "Accelerated Volume Rendering and Tomographic Reconstruction Using Texture Mapping Hardware," *Proc. 1994 Symp. Volume Visualization*, pp. 91-98, 1994.
- [2] A. Cumani, "Edge Detection in Multispectral Images," *CVGIP: Graphical Models and Image Processing*, vol. 53, pp. 40-51, 1991.
- [3] R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics (SIGGRAPH '88 Proc.)*, no. 22, pp. 65-74, 1988.
- [4] D.S. Ebert, T. McClanahan, P. Rheingans, and T. Yoo, "Direct Volume Rendering from Photographic Data," *Proc. Data Visualization 2000 (Vissym '00: Joint IEEE/Eurographics Symp. Visualization)*, pp. 137-147, May 2000.
- [5] D.S. Ebert and R.E. Parent, "Rendering and Animation of Gaseous Phenomena by Combining Fast Volume and Scanline A-Buffer Techniques," *Computer Graphics (SIGGRAPH '90 Proc.)*, F. Baskett, ed., vol. 24, pp. 357-366, Aug. 1990.
- [6] D.S. Ebert, C.D. Shaw, A. Zwa, and C. Starr, "Two-Handed Interactive Stereoscopic Visualization," *Proc. IEEE Visualization '96*, pp. 205-210, Oct. 1996.
- [7] M. Fairchild, *Color Appearance Models*. Addison-Wesley, 1998.
- [8] S. Fang, T. Biddlecome, and M. Tuceryan, "Image-Based Transfer Function Design for Data Exploration in Volume Visualization," *Proc. IEEE Visualization '98*, pp. 319-326, Oct. 1998.
- [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*, pp. 467-470, Oct. 1999.
- [10] G. Gerig, O. Kubler, R. Kikinis, and F.A. Jolesz, "Nonlinear Anisotropic Filtering of MRI Data," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, pp. 221-232, 1992.
- [11] R. Hall, *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, 1989.
- [12] T. Hazeldine, "Annular Slicing: Physics, Materials, Geology, Sedimentology and Medicine Make Use of Precision Annular Cutting. So Does the Semiconductor Industry," *Semiconductor*, Oct. 1997.
- [13] C. Imielinska, "Technical Challenges of 3D Visualization of Large Color Data Sets," *Proc. Second Visible Human Project Conf.*, R.A. Banvard and P. Cerveri, eds., Oct. 1998.
- [14] A. Kaufman, "Volume Visualization," *ACM Computing Surveys*, vol. 28, no. 1, pp. 165-167, 1996.
- [15] J. Kerr, "Photorealistic Volume-Rendered Anatomical Atlases and Interactive Virtual Dissections of the Dissectible Human," *Proc. Visible Human Project Conf.*, R.A. Banvard, ed., Oct. 1996.
- [16] G. Kindlmann and J.W. Durkin, "Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering," *Proc. 1998 Volume Visualization Symp.*, pp. 79-86, Oct. 1998.
- [17] D.H. Laidlaw, K.W. Fleischer, and A.H. Barr, "Partial-Volume Bayesian Classification of Material Mixtures in MR Volume Data Using Voxel Histograms," *IEEE Trans. Image Processing*, vol. 17, no. 1, pp. 74-86, 1998.
- [18] M. Levoy, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*, vol. 8, no. 3, pp. 29-37, 1988.
- [19] W.E. Lorensen, "Marching through the Visible Man," *Proc. IEEE Visualization '95*, pp. 368-373, Oct. 1995.
- [20] W.E. Lorensen and H.E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," *Computer Graphics (SIGGRAPH '87 Proc.)*, vol. 21, pp. 163-169, 1987.
- [21] J. Marquez, "Radiometric Inhomogeneities in the Color Cryosection Images of the VHP," *Proc. Visible Human Project Conf.*, R.A. Banvard, ed., Oct. 1996.
- [22] W. Nip and C. Logan, "Whole Frog Technical Report," Technical Report LBL-35331, Univ. of California, Lawrence Berkeley Laboratory, 1991.
- [23] H. Pfister, B. Lorensen, C. Bajaj, G. Kindlmann, W. Schroeder, L.S. Avila, K. Martin, R. Machiraju, and J. Lee, "The Transfer Function Bake-Off," *IEEE Computer Graphics and Applications*, vol. 21, no. 3, pp. 16-22, May/June 2001.
- [24] G. Sapiro and D. Ringach, "Anisotropic Diffusion of Multivalued Images with Applications to Color Filtering," *IEEE Trans. Image Processing*, vol. 5, pp. 1582-1586, 1996.
- [25] G. Sapiro, "Color Snakes," Technical Report HPL-95-113, Hewlett-Packard Labs, 1995.
- [26] G. Sapiro, *Geometric Partial Differential Equations and Image Analysis*. Cambridge Univ. Press, 2001.
- [27] S. Spitzer, M.J. Ackerman, A.L. Scherzinger, and D. Whitlock, "The Visible Human Male: A Technical Report," *J. Am. Medical Informatics Assoc.*, vol. 3, no. 2, pp. 118-130, 1996.
- [28] I. Takanashi, E.B. Lum, J. Meyer, K.-L. Ma, B. Hamann, and A. Olson, "Segmentation and Volume Rendering of Human Brain Cryosections," *Work In Progress Sessions (CD-ROM)*, IEEE CS, Oct. 2000.
- [29] A.W. Toga and J.C. Mazziotta, *Brain Mapping: The Methods*. Academic Press, 1996.
- [30] G. Wyszecki and W.S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*. John Wiley & Sons, 1982.
- [31] R. Yagel, D.S. Ebert, J.N. Scott, and Y. Kurzion, "Grouping Volume Renderers for Enhanced Visualization in Computational Fluid Dynamics," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 2, pp. 117-132, June 1995.
- [32] R. Zhou and E. Henderson, "Visualization of the Visible Human Anatomical Images," *Proc. Second Visible Human Project Conf.*, R.A. Banvard and P. Cerveri, eds., Oct. 1998.



David S. 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, mini-

mally 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, cochairing the IEEE Visualization '98 and '99 Papers program, serving on the ACM SIGGRAPH Executive Committee, and serving as an associate editor for the *IEEE Transactions on Visualization and Computer Graphics*. He is a member of the IEEE and the IEEE Computer Society.



Christopher J. Morris received the MS degree in computer science from the University of Maryland-Baltimore County in 2001, where his focus was in scientific visualization. Previously, he received the MS degree in mechanical engineering from Stanford University in 1998, where his interests included robotics and engineering design. He is an employee in the Visual Technology Group at the IBM T.J. Watson Research Center. His current research

interests include volume rendering, nonphotorealistic rendering, and parallel systems.



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 graphics,

dynamic and interactive representations and interfaces, and the experimental validation of visualization techniques. She is a member of the IEEE Computer Society.



Terry S. Yoo holds the AB degree in biology from Harvard University and the MS and PhD degrees in computer science from the University of North Carolina, Chapel Hill. He is a computer scientist in the Office of High Performance Computing and Communications, National Library of Medicine, National Institutes of Health, where he explores the processing and visualizing of 3D medical data, interactive 3D graphics, and computational geometry. Previously, as a

professor of radiology, he managed a visualization research program for interventional MRI with the University of Mississippi. He is a member of the IEEE Computer Society.

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