

# Illustrative Context-Preserving Exploration of Volume Data

Stefan Bruckner, Sören Grimm, Armin Kanitsar, and M. Eduard Gröller, *Member, IEEE*

**Abstract**—In volume rendering it is very difficult to simultaneously visualize interior and exterior structures while preserving clear shape cues. Highly transparent transfer functions produce cluttered images with many overlapping structures, while clipping techniques completely remove possibly important context information. In this paper we present a new model for volume rendering, inspired by techniques from illustration. It provides a means of interactively inspecting the interior of a volumetric data set in a feature-driven way which retains context information. The context-preserving volume rendering model uses a function of shading intensity, gradient magnitude, distance to the eye point, and previously accumulated opacity to selectively reduce the opacity in less important data regions. It is controlled by two user-specified parameters. This new method represents an alternative to conventional clipping techniques, shares their easy and intuitive user control, but does not suffer from the drawback of missing context information.

**Index Terms**—illustrative visualization, focus+context techniques, volume rendering.

## I. INTRODUCTION

**I**N direct volume rendering basically every single sample contributes to the final image allowing simultaneous visualization of surfaces and internal structures. However, in practice it is rather difficult and time-demanding to specify an appropriate transfer function. Due to the exponential attenuation, objects occluded by other semi-transparent objects are difficult to recognize. Furthermore, reducing opacities also results in reduced shading contributions per sample which causes a loss of shape cues, as shown in Figure 1 (a). One approach to overcome this difficulty is to use steep transfer functions, i.e., those with rapidly increasing opacities at particular value ranges of interest. This strengthens depth cues by creating the appearance of surfaces within the volume,

but it does so by hiding all information in some regions of the volume, sacrificing a key advantage of volume rendering. Figure 1 (b) shows an example.

Volume clipping usually plays a decisive role in understanding 3D volumetric data sets. It allows us to cut away selected parts of the volume, based on the position of voxels in the data set. Very often, clipping is the only way to uncover important, otherwise hidden, details of a data set. These clipping operations are generally not data-aware and do not take into account features of the volume. Consequently, clipping can also remove important context information leading to confusing and partly misleading result images as displayed in Figure 1 (c).

In order to resolve these issues, we propose to only suppress regions which do not contain strong features when browsing through the volume. Our idea is based on the observation that large highly lit regions usually correspond to rather homogenous areas which do not contain characteristic features. While the *position* and *shape* of specular highlights, for example, provide good cues for perceiving the curvature of a surface, the area *inside* the highlight could also be used to display other information. Thus, we propose to make this area transparent allowing the user to see the interior of the volume.

In illustrations artists use *ghosting* to visualize important internal as well as external structures. Less significant items are just faintly indicated through outlines or transparency. For example, rather flat surfaces are faded from opaque to transparent to reveal the interior of an object. Detailed structures are still displayed with high opacity. An example is shown in the technical illustration in Figure 2: while the door is almost fully transparent, small features such as the door knob are depicted opaque. The goal of this illustration technique is to provide enough hints to enable viewers to mentally complete the partly removed structures. Unlike *cutaways*, which completely remove occluding objects, *ghosting* is less invasive in the sense that important features remain visible even when they occlude structures of primary interest.

Using our idea of lighting-driven feature classification, we can easily mimic this artistic concept in an illustrative volume rendering model. Our approach allows context-

S. Bruckner and M. E. Gröller are with the Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria. E-mail: {bruckner | groeller}@cg.tuwien.ac.at

S. Grimm is with Biotronics 3D Limited, United Kingdom. E-mail: soeren@biotronics3d.com

A. Kanitsar is with TIANI Medgraph AG, Austria. E-mail: armin@tiani.com

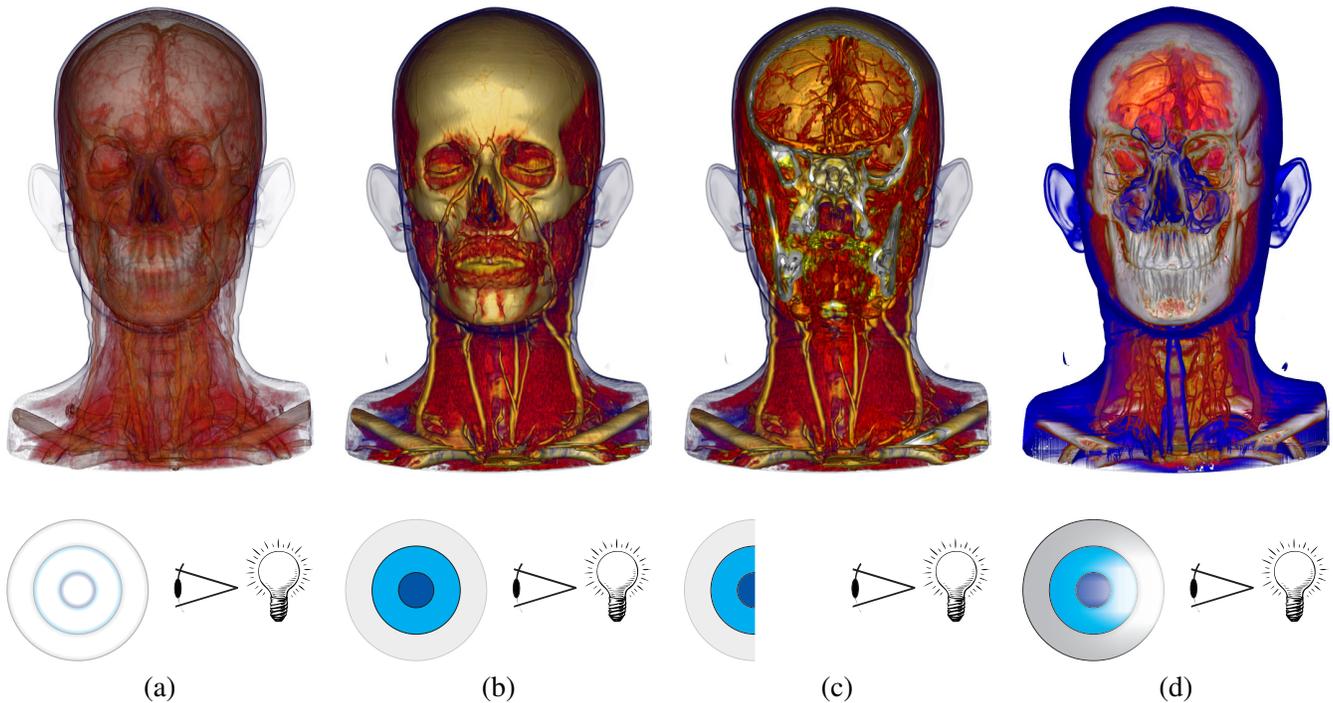


Fig. 1. Four different visualizations of a contrast-enhanced CT angiography data set. The illustration below each image depicts the basic opacity specification strategy. **(a)** Gradient-magnitude opacity-modulation. **(b)** Direct volume rendering. **(c)** Direct volume rendering with cutting plane. **(d)** Context-preserving volume rendering.

preserving and smooth clipping by adjusting the model parameters. The approach of clipping planes is extended to allow feature-aware clipping. Figure 1 (d) shows a result achieved with our method - the blood vessels inside the head are revealed while preserving context information.

This paper is based on our previous model for context-preserving exploration of volumetric data [1]. It presents a more detailed explanation of the basic model and extends it to multiple light sources. Additional applications, such as the visualization of photographic volume data and the integration into a volume segmentation pipeline, are discussed. The paper is structured as follows: Section II discusses related work. In Section III the basic context-preserving volume rendering model and its extensions are explained. Section IV presents results and applications achieved with our technique. Section V discusses the advantages of our new approach. Implementation details are given in Section VI. The paper is concluded in Section VII.

## II. RELATED WORK

We draw our inspiration mainly from three areas of research in volume visualization: volume clipping, transfer functions, and illustrative techniques.

The basic problem of clipping planes is that context information is unselectively removed when inspecting



Image courtesy of Kevin Hulsey Illustration  
Copyright ©2004 Kevin Hulsey Illustration, Inc. All rights reserved.  
<http://www.khulsey.com>

Fig. 2. Technical illustration using ghosting to display interior structures.

the interior of an object. There are approaches which try to remedy this deficiency by using more complex clipping geometry. Weiskopf et al. [2], [3] present techniques for interactively applying arbitrary convex and concave clipping objects. Konrad-Verse et al. [4] use a mesh which can be flexibly deformed by the user with an adjustable sphere of influence. Zhou et al. [5] propose the use of distance to emphasize and de-emphasize different regions. They use the distance from a focal point to directly modulate the opacity at each sample position. Thus, their approach can be seen as a generalization of binary clipping. Volume sculpting, proposed by Wang and Kaufman [6], enables interactive carving

of volumetric data. An automated way of performing view-dependent clipping operations has been presented by Viola et al. [7], [8]. Inspired by cutaway views, which are commonly used in technical illustrations, they apply different compositing strategies to prevent an object from being occluded by a less important object.

Multi-dimensional transfer functions have been proposed to extend the classification space and to allow better selection of features. Kniss et al. [9], [10] use a two-dimensional transfer function based on scalar value and gradient magnitude to effectively extract specific material boundaries and convey subtle surface properties. Hladůvka et al. [11] propose the concept of curvature-based transfer functions. Kindlmann et al. [12] employ curvature information to achieve illustrative effects, such as ridge and valley enhancement. Lum and Ma [13] assign colors and opacities as well as parameters of the illumination model through a transfer function lookup. They apply a two-dimensional transfer function to emphasize material boundaries using illumination. While there are many advantages of multi-dimensional approaches, the extension of the domain requires a well-designed user interface to cope with this increased complexity.

While the field of non-photorealistic rendering is more concerned with imitating artistic styles in an automated way, illustrative visualization goes one step further and tries to apply these techniques to enhance visual comprehension. The inherent complexity of volumetric data has led to a considerable amount of research in illustrative techniques for volume visualization. A common illustrative method is gradient-magnitude opacity-modulation. Levoy [14] proposes to modulate the opacity at a sample position using the magnitude of the local gradient. This is an effective way to enhance surfaces in volume rendering, as homogeneous regions are suppressed. Based on this idea, Ebert and Rheingans [15], [16] present several illustrative techniques which enhance features and add depth and orientation cues. They also propose to locally apply these methods for regional enhancement. Using similar methods, Lu et al. [17], [18] developed an interactive direct volume illustration system that simulates traditional stipple drawing. Csébfalvi et al. [19] visualize object contours based on the magnitude of local gradients as well as on the angle between viewing direction and gradient vector using depth-shaded maximum intensity projection. The concept of two-level volume rendering, proposed by Hauser et al. [20], [21], allows focus+context visualization of volume data. Different rendering styles, such as direct volume rendering and maximum intensity projection, are used to emphasize objects of interest while still displaying the remaining data as context.

The main contribution of this paper is a new illustrative volume rendering model which incorporates the functionality of clipping planes in a feature-aware manner. It includes concepts from artistic illustration to enhance the information content of the image. Cumbersome transfer function specification is simplified and no segmentation is required as features are classified implicitly. Our approach is especially well-suited for interactive exploration.

### III. THE CONTEXT-PRESERVING VOLUME RENDERING MODEL

Lighting plays a decisive role in illustrating surfaces. In particular, lighting variations provide visual cues regarding surface orientation. This is exploited by Lum and Ma [13], who highlight material boundaries by using two-dimensional lighting transfer functions. In contrast, in our approach the lighting intensity serves as an input to a function for opacity variation, i.e., we use the result of the shading intensity function to classify features.

#### A. Background

We assume a continuous volumetric scalar field  $f(\mathbf{P}_i)$ . A sample at position  $\mathbf{P}_i$  is denoted by  $f_{\mathbf{P}_i}$ . We denote the gradient at position  $\mathbf{P}_i$  by  $\mathbf{g}_{\mathbf{P}_i} = \nabla f(\mathbf{P}_i)$ . We use  $\hat{\mathbf{g}}_{\mathbf{P}_i}$  for the normalized gradient and  $\|\hat{\mathbf{g}}_{\mathbf{P}_i}\|_{[0..1]}$  for the gradient magnitude normalized to the interval  $[0..1]$ , where zero corresponds to the lowest gradient magnitude and one corresponds to the highest gradient magnitude in the data set.

A discrete approximation of the volume rendering integral along a viewing ray uses the front-to-back formulation of the over operator [22] to compute opacity  $\alpha_i$  and color  $c_i$  at each step along the ray:

$$\begin{aligned}\alpha_i &= \alpha_{i-1} + \alpha(\mathbf{P}_i) \cdot (1 - \alpha_{i-1}) \\ c_i &= c_{i-1} + c(\mathbf{P}_i) \cdot \alpha(\mathbf{P}_i) \cdot (1 - \alpha_{i-1})\end{aligned}\quad (1)$$

where  $\alpha(\mathbf{P}_i)$  and  $c(\mathbf{P}_i)$  are the opacity and color contributions at position  $\mathbf{P}_i$ ,  $\alpha_{i-1}$  and  $c_{i-1}$  are the previously accumulated values for opacity and color.

For conventional direct volume rendering using shading,  $\alpha(\mathbf{P}_i)$  and  $c(\mathbf{P}_i)$  are defined as follows:

$$\begin{aligned}\alpha(\mathbf{P}_i) &= \alpha_{tf}(f_{\mathbf{P}_i}) \\ c(\mathbf{P}_i) &= c_{tf}(f_{\mathbf{P}_i}) \cdot s(\mathbf{P}_i)\end{aligned}\quad (2)$$

where  $\alpha_{tf}$  and  $c_{tf}$  are the opacity and color transfer functions. They assign opacity and color to each scalar value in the volumetric function.  $s(\mathbf{P}_i)$  is the value of the shading intensity at the sample position.

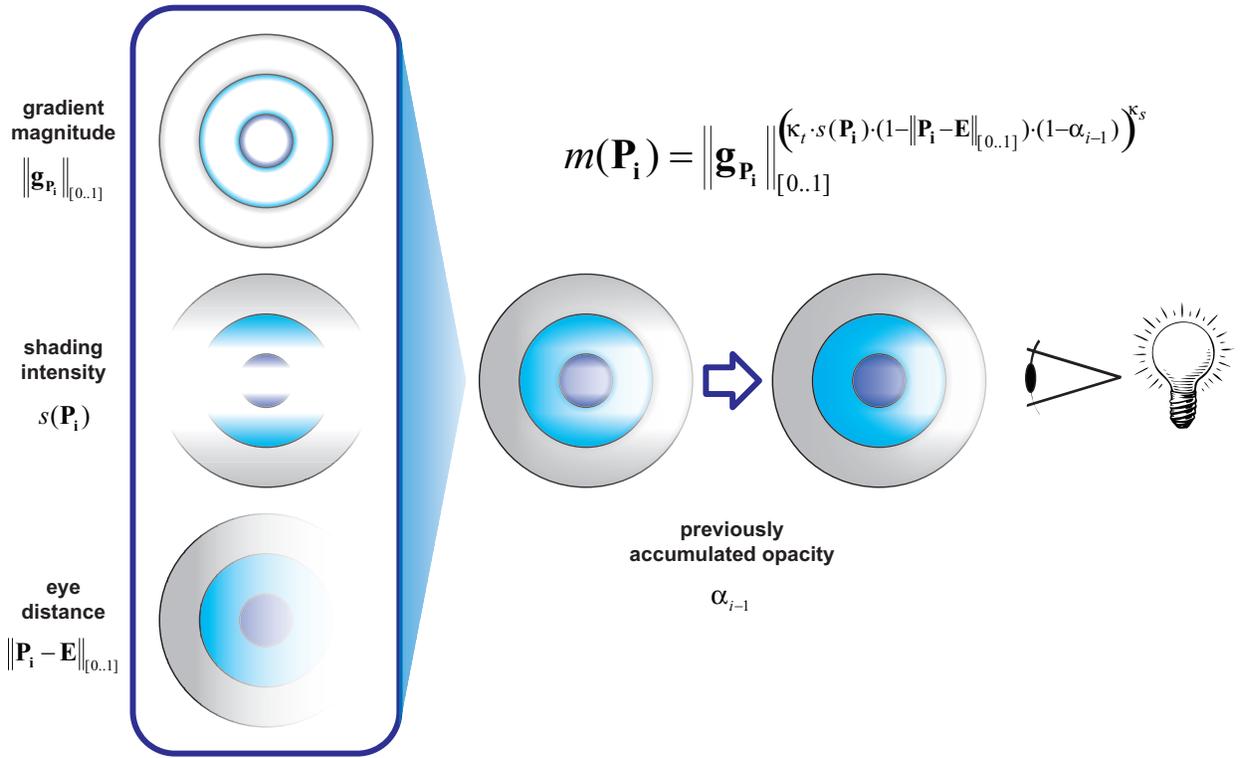


Fig. 3. Overview of the context-preserving volume rendering model. Gradient magnitude, shading intensity, and eye distance are combined to modulate sample opacity. Additionally, the effect is weighted by the previously accumulated opacity.

For the two-sided Phong-Blinn model using one directional light source,  $s(\mathbf{P}_i)$  is:

$$s(\mathbf{P}_i) = c_d \cdot \|\hat{\mathbf{L}} \cdot \hat{\mathbf{g}}_{\mathbf{P}_i}\| + c_s \cdot \left( \|\hat{\mathbf{H}} \cdot \hat{\mathbf{g}}_{\mathbf{P}_i}\| \right)^{c_e} + c_a \quad (3)$$

where  $c_d$ ,  $c_s$ , and  $c_a$  are the diffuse, specular, and ambient lighting coefficients, respectively, and  $c_e$  is the specular exponent.  $\hat{\mathbf{L}}$  is the normalized light vector and  $\hat{\mathbf{H}}$  is the normalized half-way vector. For conventional direct volume rendering, the opacity at point  $\mathbf{P}_i$  is only determined by the scalar value  $f_{\mathbf{P}_i}$  and the opacity transfer function. Gradient-magnitude opacity-modulation additionally scales the opacity by the gradient magnitude, causing an enhanced display of boundaries. For gradient-magnitude opacity-modulation  $\alpha(\mathbf{P}_i)$  is defined in the following way:

$$\alpha(\mathbf{P}_i) = \alpha_{tf}(f_{\mathbf{P}_i}) \cdot \|\mathbf{g}_{\mathbf{P}_i}\|_{[0..1]} \quad (4)$$

Within this framework,  $\alpha(\mathbf{P}_i)$  for direct volume rendering using a view-aligned clipping plane can be defined as:

$$\alpha(\mathbf{P}_i) = \begin{cases} 0 & (\mathbf{P}_i - \mathbf{E}) \cdot \hat{\mathbf{V}} < d \\ \alpha_{tf}(f_{\mathbf{P}_i}) & \text{otherwise} \end{cases} \quad (5)$$

where  $\mathbf{E}$  is the eye point,  $\hat{\mathbf{V}}$  is the normalized viewing direction and  $d$  is the depth of the clipping plane.

## B. Basic Model

Diepstraten et al. [23] cite a list of guidelines for using transparency in illustrations. For example, in many cases an illustrator chooses to reduce transparency close to the edges of transparent objects and to increase it with increasing distance to edges. Additionally, the number of overlapping transparent objects should be kept low. These rules cannot be mapped directly to the visualization of unsegmented volume data, as there is no explicit object specification. Therefore our approach relies on available quantities which are combined to achieve a similar effect.

In direct volume rendering, the illumination intensity normally does not influence the opacity at a sample position. Large regions of highly illuminated material normally correspond to rather flat surfaces that are oriented towards the light source. Our idea is to reduce the opacity in these regions. Regions, on the other hand, which receive less lighting, for example light silhouettes, will remain visible. Our model uses the result of the shading intensity function  $s(\mathbf{P}_i)$  for opacity modulation. Furthermore, to mimic - to a certain extent - the look and feel of a clipping plane, we also take the distance to the eye point into account. For viewing rays which have already accumulated a lot of opacity, we reduce subsequent opacity attenuation by our model, since too

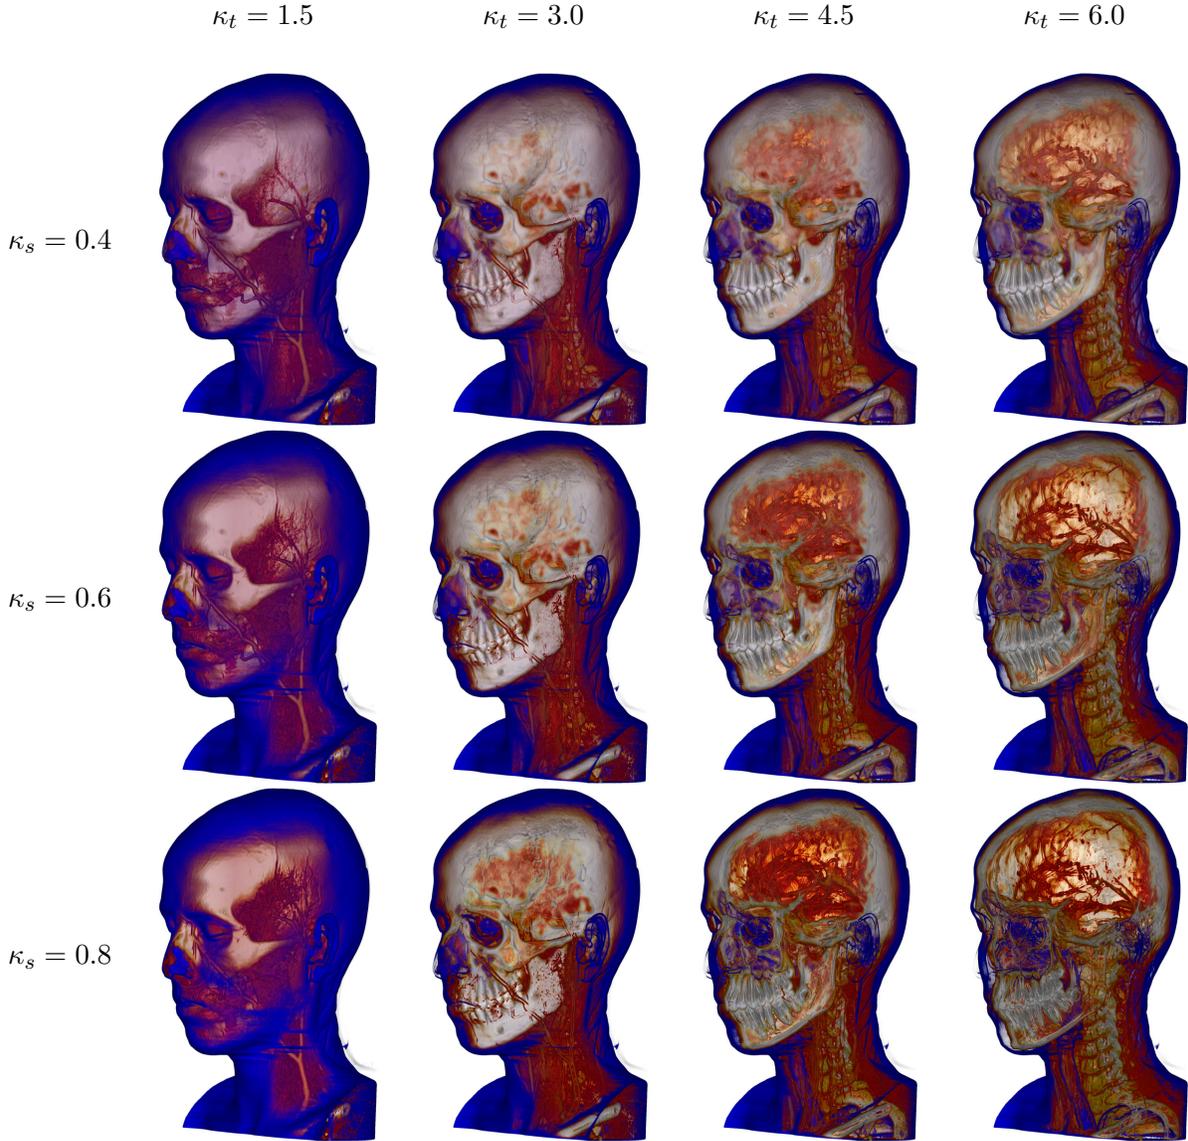


Fig. 4. Context-preserving volume rendering of a contrast-enhanced CT angiography data set using different values for  $\kappa_t$  and  $\kappa_s$ . Columns have the same  $\kappa_t$  value and rows have the same  $\kappa_s$  value.

many overlapping transparent objects make perception more difficult. As the gradient magnitude is an indicator for the homogeneity of the data, it makes sense to also take this quantity into account.

Combining available quantities in order to achieve the desired effects leads us to the following equation for the opacity at each sample position  $\mathbf{P}_i$ :

$$\alpha(\mathbf{P}_i) = \alpha_{tf}(f_{\mathbf{P}_i}) \cdot m(\mathbf{P}_i) \quad (6)$$

with

$$m(\mathbf{P}_i) = \|\mathbf{g}_{\mathbf{P}_i}\|_{[0..1]}^{(\kappa_t \cdot s(\mathbf{P}_i) \cdot (1 - \|\mathbf{P}_i - \mathbf{E}\|_{[0..1]}) \cdot (1 - \alpha_{i-1}))^{\kappa_s}} \quad (7)$$

where  $\|\mathbf{g}_{\mathbf{P}_i}\|_{[0..1]}$  is the gradient magnitude normalized to the range  $[0..1]$  and  $s(\mathbf{P}_i)$  is the shading intensity

at the current sample position. A high value of  $s(\mathbf{P}_i)$  indicates a highlight region and decreases opacity. The term  $\|\mathbf{P}_i - \mathbf{E}\|_{[0..1]}$  is the distance of the current sample position  $\mathbf{P}_i$  to the eye point  $\mathbf{E}$ , normalized to the range  $[0..1]$ . Zero corresponds to the sample position closest to the eye point and one corresponds to the sample position farthest from the eye point. Thus, the effect of our model will decrease as distance increases. Due to the term  $1 - \alpha_{i-1}$  structures located behind semi-transparent regions will appear more opaque.

Figure 3 illustrates the effect of these different quantities: regions with low gradient magnitude (left, top) and high shading intensity (left, center) are more likely to be suppressed. Due to the inclusion of the eye distance (left, bottom) this suppression will be strongest for regions close to the viewer. Additionally, this effect is dependent

on the previously accumulated opacity (right, center). For viewing rays that already have accumulated high opacity, further opacity modulation will be reduced effectively making any further samples more opaque.

The model is controlled by the two user-specified parameters  $\kappa_t$  and  $\kappa_s$ . Through modification of these parameters the user can interactively uncover occluded regions. The parameter  $\kappa_t$  roughly corresponds - due to the position-dependent term  $1 - \|\mathbf{P}_i - \mathbf{E}\|_{[0..1]}$  - to the depth of a clipping plane, i.e., higher values reveal more of the interior of the volume. The effect of modifying  $\kappa_s$  is less pronounced; it basically controls the sharpness in transition between clipped and visible regions. Higher values will result in very sharp cuts, while lower values produce smoother transitions.

Figure 4 shows results for different settings of  $\kappa_t$  and  $\kappa_s$ . As  $\kappa_t$  increases more of the interior of the head is revealed - structures on the outside are more and more reduced to the most prominent features. An increase in  $\kappa_s$  causes a sharper transition between attenuated and visible regions. An interesting property of this model is that it is a unifying extension of both direct volume rendering and gradient-magnitude opacity-modulation. If  $\kappa_t$  is set to zero, the opacity remains unmodified and normal direct volume rendering is performed. Likewise, when  $\kappa_s$  is set to zero, the opacity is directly modulated by the gradient magnitude.

There are several reasons for choosing an exponential mapping in (7). By placing the gradient magnitude in the mantissa and the modifying terms in the exponent, the highest gradient magnitudes in the data will always be preserved, as they are close to one, while lower gradient magnitudes will be more affected by the modulation. This ensures that the central role of the gradient magnitude in defining features in a data set of unknown nature is accounted for. In the exponent the term  $(\kappa_t \cdot (\dots))^{\kappa_s}$  is chosen as it allows fine-grained control over the shape of the function similar to Phong's specular term. While  $\kappa_t$  controls the basic slope,  $\kappa_s$  allows to adjust the curvature of this function.

### C. Multiple Light Sources

While useful results can be achieved with the basic model featuring only one light source, an obvious extension is the inclusion of multiple light sources. For  $N$  light sources we assume a set of illumination intensity functions  $s_0, s_1, \dots, s_{N-1}$  and the parameters  $\kappa_{s_0}, \kappa_{s_1}, \dots, \kappa_{s_{N-1}}$  and  $\kappa_{t_0}, \kappa_{t_1}, \dots, \kappa_{t_{N-1}}$ . Any of these light sources can either exhibit context-preserving opacity-modulation or act as a conventional light source. The indices of all context-preserving light sources are members of the index set  $T$ .



Fig. 5. Context-preserving volume rendering using two light sources. The context-preserving light source is placed in front of the face. A conventional light source illuminates the head from the left.

We redefine (7) to take into account all context-preserving light sources:

$$m(\mathbf{P}_i) = \prod_{j \in T} \|\mathbf{g}_{\mathbf{P}_i}\|_{[0..1]}^{(\kappa_{t_j} \cdot s_j(\mathbf{P}_i) \cdot (1 - \|\mathbf{P}_i - \mathbf{E}\|_{[0..1]}) \cdot (1 - \alpha_{i-1}))^{\kappa_{s_j}}} \quad (8)$$

Furthermore, we need to take the illumination contributions of all light sources into account for the final color at the resample point  $c(\mathbf{P}_i)$ :

$$c(\mathbf{P}_i) = c_{tf}(f_{\mathbf{P}_i}) \cdot \sum_{j=0}^{N-1} s_j(\mathbf{P}_i) \quad (9)$$

Apart from these minor modifications no other changes to the basic model are required. The advantage of using multiple light sources is that the exploration is not influenced by basic lighting settings. A typical setup might employ one conventional light source for normal illumination and one context-preserving light source used for examining the data, as shown in Figure 5.

### D. Data-Dependent Parameters

The two parameters of our model  $\kappa_t$  and  $\kappa_s$  allow intuitive control of the visualization:  $\kappa_t$  is used to interactively browse through the volume, similar to the variation

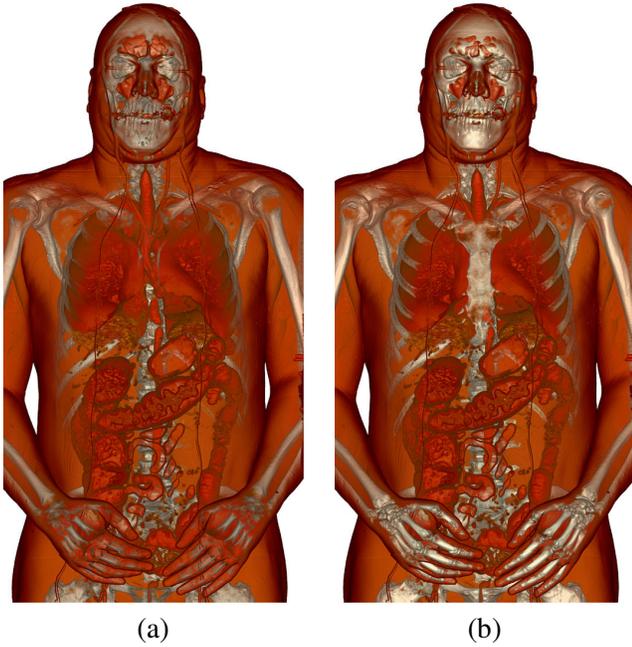


Fig. 6. Context-preserving volume rendering of the visible human male CT data set. (a) Only global parameter settings are used. (b) Bone is made impenetrable by using data-dependent parameters.

of the depth of a clipping plane;  $\kappa_s$  normally remains fixed during this process and is only later adjusted to achieve a visually pleasing result.

While we have found that these parameters provide sufficient control in most cases, a possible extension is to make them data dependent, i.e., they are defined by specifying a transfer function. This increases the flexibility of the method, but also raises the burden on the user, as transfer function specification is a complex task. Thus, we propose a hybrid solution between both approaches. We keep the global constants  $\kappa_t$  and  $\kappa_s$ , but their values are modulated by rather simple data-dependent functions. In (7),  $\kappa_t$  is replaced by  $\kappa_t \cdot \lambda_t(\mathbf{P}_i)$  and  $\kappa_s$  is replaced by  $\kappa_s \cdot \lambda_s(\mathbf{P}_i)$ . Both  $\lambda_t$  and  $\lambda_s$  are real-valued functions in the range  $[0..1]$ . For example, the user can specify zero for  $\lambda_t$  to make some regions impenetrable. Likewise, setting  $\lambda_s$  to zero for certain values ensures pure gradient-magnitude opacity-modulation. If one of these functions has a value of one everywhere, the corresponding global parameter remains unchanged. These modulation functions can be based on the scalar value, segmentation information (region identifier) or similar properties.

Figure 6 (a) shows the visible human male CT data set rendered using just the global parameters, while in Figure 6 (b) bone is made impenetrable by setting  $\lambda_t$  to zero for the corresponding values.

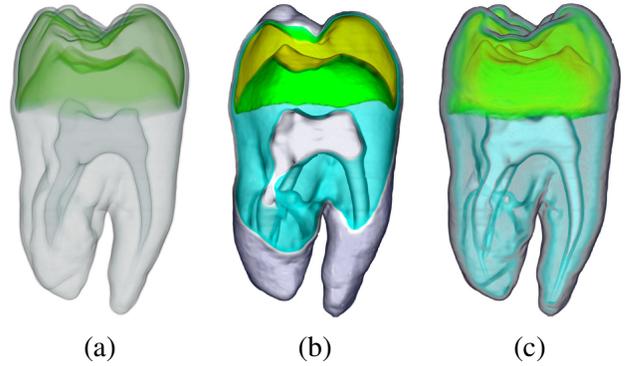


Fig. 7. CT scan of a tooth rendered with three different techniques. (a) Gradient-magnitude opacity-modulation. (b) Direct volume rendering with one clipping plane. (c) Context-preserving volume rendering.

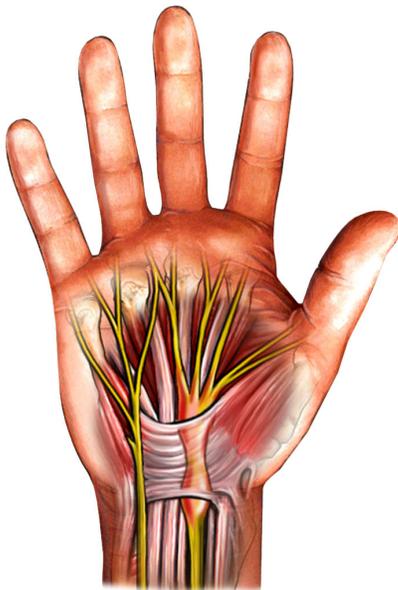
#### IV. RESULTS AND APPLICATIONS

We experimented with the presented model using a wide variety of volumetric data sets. We have found that our approach makes transfer function specification much easier, as there is no need to pay special attention to opacity. Normally, tedious tuning is required to set the appropriate opacity in order to provide good visibility for all structures of interest. Using the context-preserving volume rendering model, we just assign colors to the structures and use the parameters  $\kappa_t$  and  $\kappa_s$  to achieve an insightful rendering. Opacity in the transfer function is just used to suppress certain regions, such as background. This contrasts the usual direct volume rendering approach, where opacity specification is vital in order to achieve the desired visual result. In many cases, however, good results are difficult and laborious to achieve. For example, for structures sharing the same value range, as it is often the case with contrast-enhanced CT scans, it is impossible to assign different opacities using a one-dimensional transfer function. If one object is occluding the other, setting a high opacity will cause the occluded object to be completely hidden. Using high transparency, on the other hand, will make both objects hardly recognizable. Our method inherently solves this issue, as it bases opacity not only on data values, but also includes a location-dependent term. In the following, we present some results achieved with our model in combination with a one-dimensional color transfer function. No segmentation was applied.

Figure 7 shows the tooth data set rendered with gradient-magnitude opacity-modulation, direct volume rendering using a clipping plane, and context-preserving volume rendering using the same transfer function. Gradient-magnitude opacity-modulation shows the whole data set but the overlapping transparent structures make the interpretation of the image a difficult task. On



(a)



(b)

Image courtesy of Nucleus Medical Art  
 Copyright ©2004 Nucleus Medical Art, Inc. All rights reserved.  
<http://www.nucleusinc.com>

Fig. 8. Comparing context-preserving volume rendering to illustration. (a) Context-preserving volume rendering of a hand data set. (b) Medical illustration using ghosting.

the other hand, it is very difficult to place a clipping plane in a way that it does not remove features of interest. Using context-preserving volume rendering, the clipping depth adapts to features characterized by varying lighting intensity or high gradient magnitude.

Figure 8 (a) shows a CT scan of a human hand rendered using our model. The resulting image has a strong resemblance to medical illustrations using the *ghosting* technique, as can be seen by comparing Figure 8 (a) and (b). By preserving certain characteristic features, such as

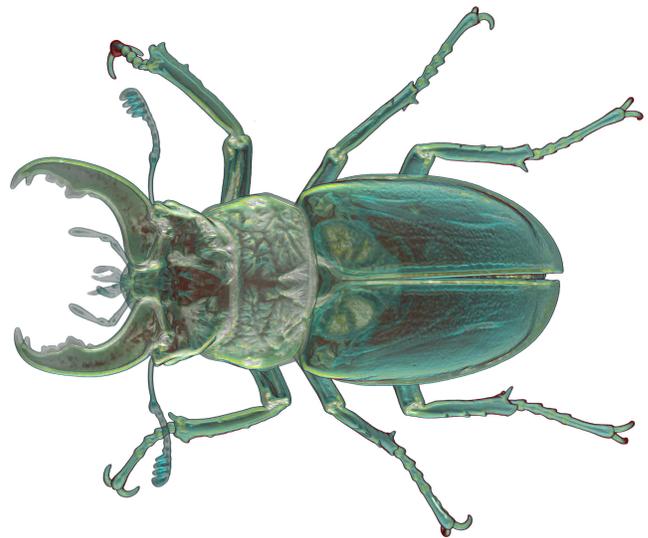


Fig. 9. Context-preserving volume rendering of a CT scan of a stag beetle.

creases on the skin, and gradually fading from opaque to transparent, the human mind is able to reconstruct the whole object from just a few hints while inspecting the detailed internal structures. This fact is commonly exploited by illustrators for static images. For interactive exploration the effect becomes even more pronounced and causes a very strong impression of depth.

In Figure 9 a context-preserving volume rendering of a CT scan of a stag beetle is depicted. Highly detailed internal structures can be clearly identified while the translucent exterior provides the necessary context. The shading-dependent term in our model causes reduced transparency near the edges which serves as a cue for the relative depth of different structures.

In Figure 10 context-preserving volume rendering is applied to a contrast-enhanced CT scan of a human body. Two light sources are used: a conventional light source is illuminating the body from the front and a context-preserving light source is used to reveal the interior from the side. Our method allows to create a highly detailed ghosted view of the human body without any segmentation.

A typical application of traditional clipping planes is volume rendering of stacks of histological cross-sections. For such data sets it is not easily possible to assign different opacities to particular structures using transfer functions. Often it is only possible to separate the actual object from its background. Without prior segmentation a common strategy is to display the entire volume employing clipping planes to explore its interior. As the context-preserving volume rendering model uses relative local quantities to compute the opacity of a sample, it does not suffer from the same problems as

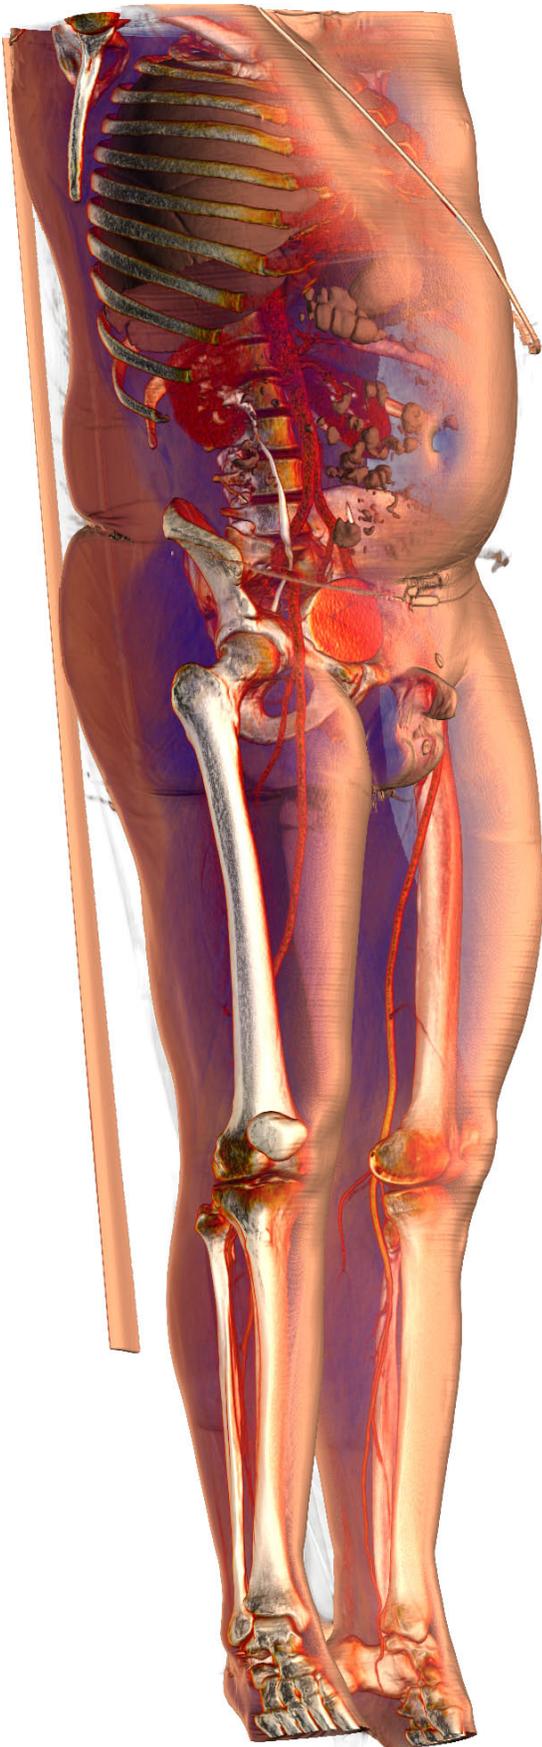


Fig. 10. Context-preserving volume rendering of a full-body CT angiography scan.

global transfer functions which assign opacities based on absolute values. The scalar gradient is replaced by an appropriate metric for photographic data, such as the color distance gradient proposed by Ebert et al. [24]. Figure 11 shows that our model allows exploration of color volumes without prior segmentation. Different structures are accentuated as the model parameters are modified. This approach might not completely alleviate the need for segmentation in some applications, but it can act as a tool for exploring and previewing complex spatial relationships as part of a segmentation pipeline.

One usage scenario for the presented visualization technique is an interactive segmentation application for photographic data. The unsegmented data can be visualized three-dimensionally as the segmentation is in progress. This supports the labeling of complex structures which is often difficult based on the original slices. Already segmented features can be made impenetrable (see Section III-D), while the rest of the data is displayed using context-preserving volume rendering. An example is shown in Figure 12. With our method effective volume visualizations of partly segmented data can be generated easily. Visual properties of already segmented objects can be controlled as usual, while visibility of the remaining data is determined through the context-preserving rendering model.

As some of the strengths of our model are most visible in animated viewing, several supplementary video sequences are available for download at: <http://www.cg.tuwien.ac.at/research/publications/2006/bruckner-2006-ICE>

## V. DISCUSSION

Context-preserving volume rendering presents an alternative to conventional clipping techniques. It provides a simple interface for examining the interior of volumetric data sets. In particular, it is well-suited for medical data which commonly have a layered structure. Our method provides a mechanism to investigate structures of interest that are located inside a larger object with similar value ranges, as it is often the case with contrast-enhanced CT data. Landmark features of the data set are preserved. Our approach does not require any form of pre-processing, such as segmentation. One key feature of this model is that it effectively simplifies opacity definition during the transfer function specification.

Context-preserving volume rendering does not pose any restrictions on the type of transfer function used. The model could be applied without changes to modulate the opacity retrieved from a multi-dimensional transfer function. Likewise, the modulation functions  $\lambda_t$  and  $\lambda_s$

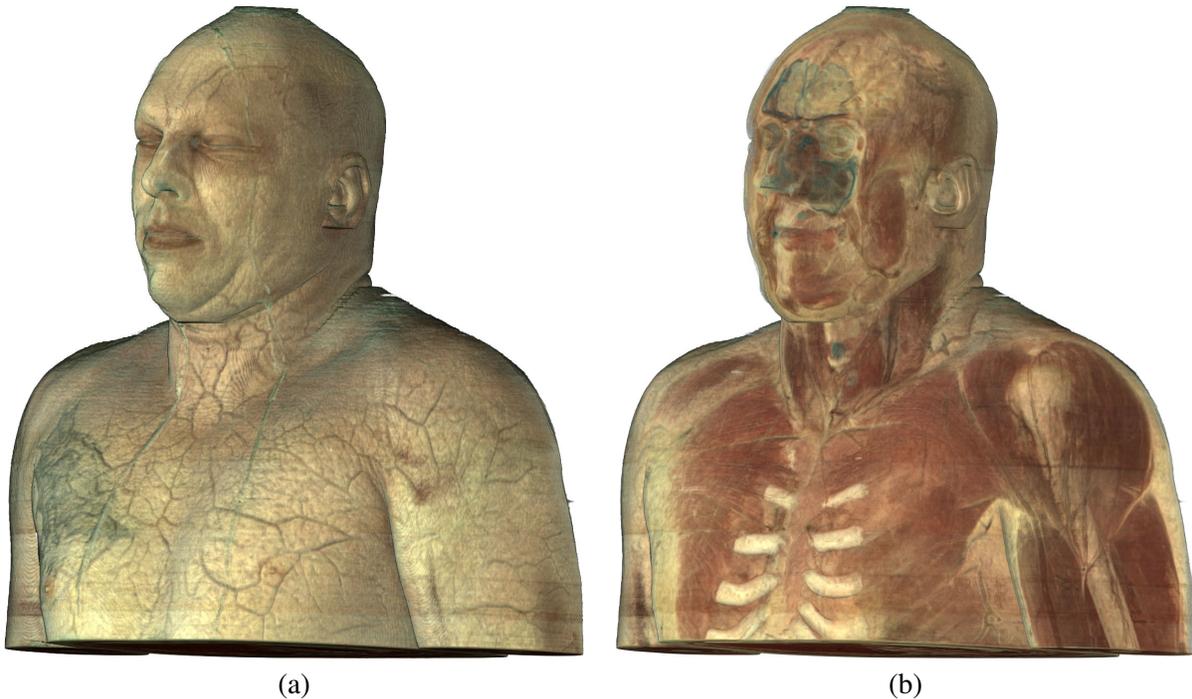


Fig. 11. Using context-preserving volume rendering for exploring photographic data. **(a)** A low value of  $\kappa_t$  reveals superficial blood vessels. **(b)** A higher value of  $\kappa_t$  displays muscle and brain tissue.

could be multi-dimensional. For reasons of simplicity, we have only considered simple one-dimensional transfer functions in this paper.

Further degrees of freedom of the method are provided by its close connection to the illumination model. By changing the direction of a directional light source, for example, features can be interactively highlighted or suppressed, based on their orientation. Modifying the diffuse and specular factors will result in a variation of directional dependency, while adjusting the ambient component has a global influence. As means of changing these illumination properties are included in every volume visualization system, this additional flexibility will not increase the complexity of the user interface.

## VI. IMPLEMENTATION

The implementation of context-preserving volume rendering is straight-forward, as it only requires a simple addition in the compositing routine of an existing volume rendering algorithm. The model uses only quantities which are commonly available in every volume renderer, such as gradient direction and magnitude and the depth along a viewing ray. We have integrated this method into a high-quality software volume ray casting system for large data [25], [26]. In our implementation, the most costly parts of the context-preserving volume rendering model are the exponentiations. As with the specular term of the Phong illumination model, it is sufficient

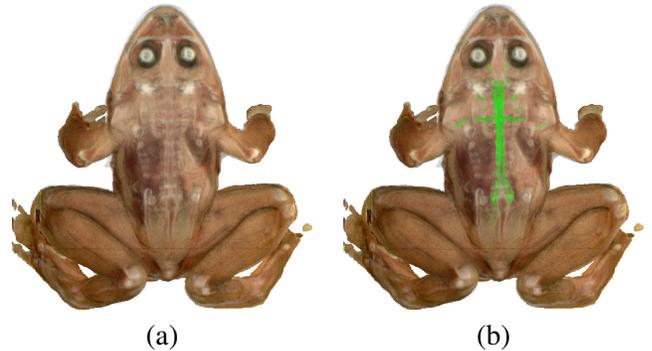


Fig. 12. Context-preserving volume rendering of the Whole Frog Project photographic data set. **(a)** Context-preserving volume rendering of the unsegmented data. **(b)** Nerves have been segmented, are made impenetrable, and are artificially colored.

to approximate the exponentiation with a function that evokes a similar visual impression. Schlick [27] proposes a simple approximation:

$$x^n \approx \frac{x}{n - nx + x} \quad (10)$$

We can approximate the two exponentiations used by our model in the following way:

$$(a \cdot y)^b \approx \frac{x \cdot (b + ay - aby)}{ay + b \cdot (x - axy)} \quad (11)$$

The approximation shares important characteristics with the original function, i.e., it is 0 for  $x = 0$ , 1 for  $y = 0$ , 1 for  $a = 0$ , and  $x$  for  $b = 0$ . Thus, the basic

behavior of our model parameters remains unchanged and the approximation error can be compensated by slightly modifying the parameter values to evoke an almost identical visual result. The advantage is that (11) can be efficiently implemented using just 4 multiplications, 2 additions, 2 subtractions, and 1 division. In our software implementation, this optimization reduced the cost from a 15% to a 5% increase compared to normal direct volume rendering.

Implementing context-preserving volume rendering on current GPUs is also straight-forward by using a single-pass raycasting fragment program [28]. We implemented a GLSL shader which is equivalent in functionality to the CPU-based version. In our tests on an NVidia 6800 GT graphics card it was not necessary to use the approximation above as exponentiation is heavily optimized due to its frequent use in shading. Typical frame rates for common test data sets such as the UNC head ( $256 \times 256 \times 224$ ) are around 10 frames per second (depending on the degree of transparency) for a  $512 \times 512$  viewport and an object sample distance of 1.0.

## VII. CONCLUSION

The focus of our research is to develop an effective alternative to conventional clipping techniques for explorative volume visualization. The key feature of our approach is that context information is preserved when inspecting the interior of a data set without requiring prior segmentation.

The context-preserving volume rendering model is inspired by ghosted views used in illustration. The interior of highlighted regions is made transparent to illustrate structures behind. The opacity of a sample is modulated by a function of shading intensity, gradient magnitude, distance to the eye point, and previously accumulated opacity. The model is controlled by two parameters which provide a simple user-interface and have intuitive meaning. Context-preserving volume rendering is a unifying extension of both direct volume rendering and gradient-magnitude opacity-modulation and allows a smooth transition between these techniques.

The approach simplifies transfer function specification, as the user only needs to specify constant opacity for structures of interest. Our method is applicable to data domains like photographic cross-sections where opacity specification is particularly difficult. Variation of the model parameters can be used to interactively explore the volumetric data set without prior knowledge of contained structures. The achieved results have a strong resemblance to artistic illustrations. Despite their increased information content, they are easy to interpret.

The approach adds little complexity to conventional direct volume rendering and it is well-suited for interactive viewing and exploration.

## ACKNOWLEDGMENT

The work presented in this publication is carried out as part of the **exvisation** project (<http://www.cg.tuwien.ac.at/research/vis/exvisation>) supported by the Austrian Science Fund (FWF) grant no. P18322.

The stag beetle data set has been provided by Georg Glaeser, Vienna University of Applied Arts, Austria and Johannes Kastner, Wels College of Engineering, Austria. The visible human data set is courtesy of the Visible Human Project [29], National Library of Medicine, USA. The frog data set is courtesy of the Whole Frog Project [30], Lawrence Berkeley National Laboratory, USA.

## REFERENCES

- [1] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller, "Illustrative context-preserving volume rendering," in *Proceedings of EuroVis 2005*, 2005, pp. 69–76.
- [2] D. Weiskopf, K. Engel, and T. Ertl, "Volume clipping via per-fragment operations in texture-based volume visualization," in *Proceedings of IEEE Visualization 2002*, 2002, pp. 93–100.
- [3] —, "Interactive clipping techniques for texture-based volume visualization and volume shading," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, no. 3, pp. 298–312, 2003.
- [4] O. Konrad-Verse, B. Preim, and A. Littmann, "Virtual resection with a deformable cutting plane," in *Proceedings of Simulation und Visualisierung 2004*, 2004, pp. 203–214.
- [5] J. Zhou, A. Döring, and K. D. Tönnies, "Distance based enhancement for focal region based volume rendering," in *Proceedings of Bildverarbeitung für die Medizin 2004*, 2004, pp. 199–203.
- [6] S. W. Wang and A. E. Kaufman, "Volume sculpting," in *Proceedings of the Symposium on Interactive 3D Graphics 1995*, 1995, pp. 151–156.
- [7] I. Viola, A. Kanitsar, and M. E. Gröller, "Importance-driven volume rendering," in *Proceedings of IEEE Visualization 2004*, 2004, pp. 139–145.
- [8] —, "Importance-driven feature enhancement in volume visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 4, pp. 408–418, 2005.
- [9] J. Kniss, G. Kindlmann, and C. Hansen, "Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets," in *Proceedings of IEEE Visualization 2001*, 2001, pp. 255–262.
- [10] —, "Multidimensional transfer functions for interactive volume rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 3, pp. 270–285, 2002.
- [11] J. Hladůvka, A. König, and M. E. Gröller, "Curvature-based transfer functions for direct volume rendering," in *Proceedings of the Spring Conference on Computer Graphics 2000*, 2000, pp. 58–65.

- [12] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Möller, "Curvature-based transfer functions for direct volume rendering: Methods and applications," in *Proceedings of IEEE Visualization 2003*, 2003, pp. 513–520.
- [13] E. B. Lum and K.-L. Ma, "Lighting transfer functions using gradient aligned sampling," in *Proceedings of IEEE Visualization 2004*, 2004, pp. 289–296.
- [14] M. Levoy, "Display of surfaces from volume data," *IEEE Computer Graphics and Applications*, vol. 8, no. 3, pp. 29–37, 1988.
- [15] D. S. Ebert and P. Rheingans, "Volume illustration: non-photorealistic rendering of volume models," in *Proceedings of IEEE Visualization 2000*, 2000, pp. 195–202.
- [16] P. Rheingans and D. S. Ebert, "Volume illustration: Nonphotorealistic rendering of volume models," *IEEE Transactions on Visualization and Computer Graphics*, vol. 7, no. 3, pp. 253–264, 2001.
- [17] A. Lu, C. J. Morris, D. S. Ebert, P. Rheingans, and C. Hansen, "Non-photorealistic volume rendering using stippling techniques," in *Proceedings of IEEE Visualization 2002*, 2002, pp. 211–218.
- [18] A. Lu, C. J. Morris, J. Taylor, D. S. Ebert, C. Hansen, P. Rheingans, and M. Hartner, "Illustrative interactive stipple rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, no. 2, pp. 127–138, 2003.
- [19] B. Csébfalvi, L. Mroz, H. Hauser, A. König, and M. E. Gröller, "Fast visualization of object contours by non-photorealistic volume rendering," *Computer Graphics Forum*, vol. 20, no. 3, pp. 452–460, 2001.
- [20] H. Hauser, L. Mroz, G.-I. Bisch, and M. E. Gröller, "Two-level volume rendering - fusing MIP and DVR," in *Proceedings of IEEE Visualization 2000*, 2000, pp. 211–218.
- [21] —, "Two-level volume rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 7, no. 3, pp. 242–252, 2001.
- [22] T. Porter and T. Duff, "Compositing digital images," *Computer Graphics*, vol. 18, no. 3, pp. 253–259, 1984.
- [23] J. Diepstraten, D. Weiskopf, and T. Ertl, "Transparency in interactive technical illustrations," *Computer Graphics Forum*, vol. 21, no. 3, 2002.
- [24] D. S. Ebert, C. J. Morris, P. Rheingans, and T. S. Yoo, "Designing effective transfer functions for volume rendering from photographic volumes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 2, pp. 183–197, 2002.
- [25] S. Grimm, S. Bruckner, A. Kanitsar, and M. E. Gröller, "Memory efficient acceleration structures and techniques for CPU-based volume raycasting of large data," in *Proceedings of the IEEE/SIGGRAPH Symposium on Volume Visualization and Graphics 2004*, 2004, pp. 1–8.
- [26] —, "A refined data addressing and processing scheme to accelerate volume raycasting," *Computers & Graphics*, vol. 28, no. 5, pp. 719–729, 2004.
- [27] C. Schlick, "A fast alternative to Phong's specular model," in *Graphics gems IV*, P. Heckbert, Ed. Academic Press, 1994, pp. 385–387.
- [28] S. Stegmaier, M. Strengert, T. Klein, and T. Ertl, "A simple and flexible volume rendering framework for graphics-hardware-based raycasting," in *Proceedings of the International Workshop on Volume Graphics 2005*, 2005, pp. 187–195.
- [29] S. Spitzer, M. J. Ackermann, A. L. Scherzinger, and D. Whitlock, "The visible human male: A technical report," *Journal of the American Medical Informatics Association*, vol. 3, no. 2, pp. 118–139, 1996.
- [30] W. Nip and C. Logan, "Whole frog technical report," University of California, Lawrence Berkeley Laboratory, Tech. Rep. LBL-35331, 1991.



**Stefan Bruckner** is a research assistant and PhD candidate at the Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria. He received his master's degree in 2004 from the same university. His research interests include illustrative visualization, medical visualization, and volume rendering.



**Sören Grimm** is director of research and development at Biotronic 3D Limited, United Kingdom. He received his PhD in 2005 from the Vienna University of Technology, Austria. His research is focused on medical visualization and volume rendering.



**Armin Kanitsar** is key developer and researcher in the 3D department of TIANI Medgraph AG, Austria. He received his PhD in 2004 from the Vienna University of Technology, Austria. His main research interests include volume visualization, medical visualization, and image processing.



**M. Eduard Gröller** is associate professor at the Vienna University of Technology. In 1993 he received his PhD from the same university. He is head of the visualization group at the Institute of Computer Graphics and Algorithms. His research interests include computer graphics, volume visualization, and medical visualization.