EUROGRAPHICS - IEEE VGTC Symposium on Visualization (2005)
K. W. Brodlie, D. J. Duke, K. I. Joy (Editors)

Illustrative Context-Preserving Volume Rendering

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

{bruckner | grimm | kanitsar | groeller}@cg.tuwien.ac.at
Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

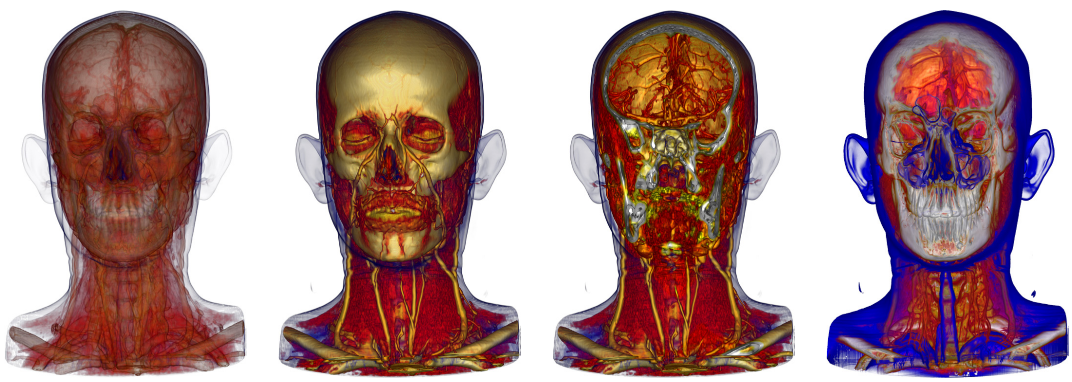

(a) (b) (c) (d)

Figure 1: *Contrast-enhanced CT angiography data set.* ***(a)*** *Gradient-magnitude opacity-modulation.* ***(b)*** *Direct volume rendering.* ***(c)*** *Direct volume rendering with cutting plane.* ***(d)*** *Context-preserving volume rendering.*

Abstract

In volume rendering it is very difficult to simultaneously visualize interior and exterior structures while preserving clear shape cues. Very 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 that 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.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Theoretically, in direct volume rendering 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 opac-

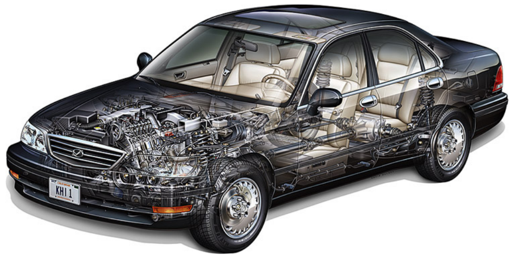

Figure 2: *Technical illustration using ghosting to display interior structures.*

ities at particular value ranges of interest. This may increase 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.

Thus, 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. However, as these clipping operations are generally not data-aware they 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 regions of high lighting intensity usually correspond to rather homogenous areas which do not contain characteristic features. While the *position* and *shape* of specular highlights, for example, give 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 illustration, when artists want to visualize specific internal structures as well as the exterior, the use of *ghosting* is common where less significant items are reduced to an illusion of 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, as shown in the technical illustration in Figure 2. The goal of this illustration technique is to provide enough hints to enable viewers to mentally complete the partly removed structures.

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-preserving 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 out method - the blood vessels inside the head are revealed while preserving context information.

2. Related Work

We draw our inspiration mainly from three areas of research in volume visualization: volume clipping, illustrative volume rendering, and transfer functions. The term illustrative volume rendering refers to approaches dealing with feature enhancement, alternative shading and compositing models, and related matters, which are usually called non-photorealistic volume rendering methods. We believe that the adjective illustrative more accurately describes the nature of many of these techniques.

Volume Clipping. There are approaches which try to remedy the deficiencies of simple clipping planes by using more complex clipping geometry. Weiskopf et al. [WEE02, WEE03] presented techniques for interactively applying arbitrary convex and concave clipping objects. Konrad-Verse et al. [KVPL04] use a mesh which can be flexibly deformed by the user with an adjustable sphere of influence. Zhou et al. [ZDT04] propose the use of distance to emphasize and de-emphasize different regions. They use the distance from the eye 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 [WK95], enables interactive carving of volumetric data. An automated way of performing clipping operations has been presented by Viola et al. [VKG04]. 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.

Illustrative Volume Rendering. A common illustrative method is gradient-magnitude opacity-modulation. Levoy [Lev88] proposed 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 [ER00, RE01] presented 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. [LME*02] developed an interactive direct volume illustration system that simulates traditional stipple drawing. Csébfalvi et al. [CMH*01] 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. [HMBG01], 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.

Transfer Functions. Multi-dimensional transfer functions have been proposed to extend the classification space and thus allow better selection of features. These transfer functions take into account derivative information. For example, Kindlmann et al. [KWTM03] use curvature information to achieve illustrative effects, such as ridge and valley enhancement. Lum and Ma [LM04] assign colors and opacities as well as parameters of the illumination model through a transfer function look up. They apply a two-dimensional transfer function to emphasize material boundaries using illumination. While there are many advantages of multi-dimensional approaches, they also have some issues: it is, for example, difficult to develop a simple and intuitive user interface which allows the specification of multi-dimensional transfer functions. One possible solution was presented by Kniss et al. [KKH01] who introduce probing and classification widgets.

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. Thus, our approach is especially well-suited for interactive exploration.

3. The Context-Preserving Volume Rendering Model

Lighting plays a critically important role in illustrating surfaces. In particular, lighting variations provide visual cues regarding surface orientation. This is acknowledged by Lum and Ma [LM04], 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 which varies the opacity based on this information, i.e., we use the result of the shading intensity function to classify features.

3.1. Background

We assume a continuous volumetric scalar field $f(P_i)$. A sample at position P_i is denoted by f_{P_i}. We denote the gradient at position P_i by $g_{P_i} = \nabla f(P_i)$. We use $\hat{g}_{P_i}$ for the normalized gradient and $\|g_{P_i}\|$ 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 to compute the opacity α_i and the color c_i at each step along the ray:

$$\begin{aligned} \alpha_i &= \alpha_{i-1} + \alpha(P_i) \cdot (1 - \alpha_{i-1}) \\ c_i &= c_{i-1} + c(P_i) \cdot \alpha(P_i) \cdot (1 - \alpha_{i-1}) \end{aligned} \tag{1}$$

$\alpha(P_i)$ and $c(P_i)$ are the opacity and color contributions at position P_i, α_{i-1} and c_{i-1} are the previously accumulated values for opacity and color.

For conventional direct volume rendering using shading, $\alpha(P_i)$ and $c(P_i)$ are defined as follows:

$$\begin{aligned} \alpha(P_i) &= \alpha_{tf}(f_{P_i}) \\ c(P_i) &= c_{tf}(f_{P_i}) \cdot s(P_i) \end{aligned} \tag{2}$$

α_{tf} and c_{tf} are the opacity and color transfer functions; they map an opacity and color to each scalar value in the volumetric function. $s(P_i)$ is the value of the shading intensity at the sample position.

For the Phong-Blinn model using one directional light source, $s(P_i)$ is:

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

c_d, c_s, and c_a are the diffuse, specular, and ambient lighting coefficients, respectively, and c_e is the specular exponent. $\hat{L}$ is the normalized light vector and $\hat{H}$ is the normalized half-way vector. For conventional direct volume rendering, the opacity at point P_i is only determined by the scalar value f_{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. Thus, for gradient-magnitude opacity-modulation $\alpha(P_i)$ is defined in the following way:

$$\alpha(P_i) = \alpha_{tf}(f_{P_i}) \cdot \|g_{P_i}\| \tag{4}$$

3.2. Model

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 contours, will remain visible.

Therefore we choose to use the result of the shading intensity function $s(P_i)$ for opacity modulation. Furthermore, as we want to mimic - to a certain extent - the look and feel of a clipping plane, we also want to take the distance to the eye point into account. For viewing rays having already accumulated a lot of opacity, we want to reduce the attenuation by our model. The total effect should also take the gradient magnitude into account.

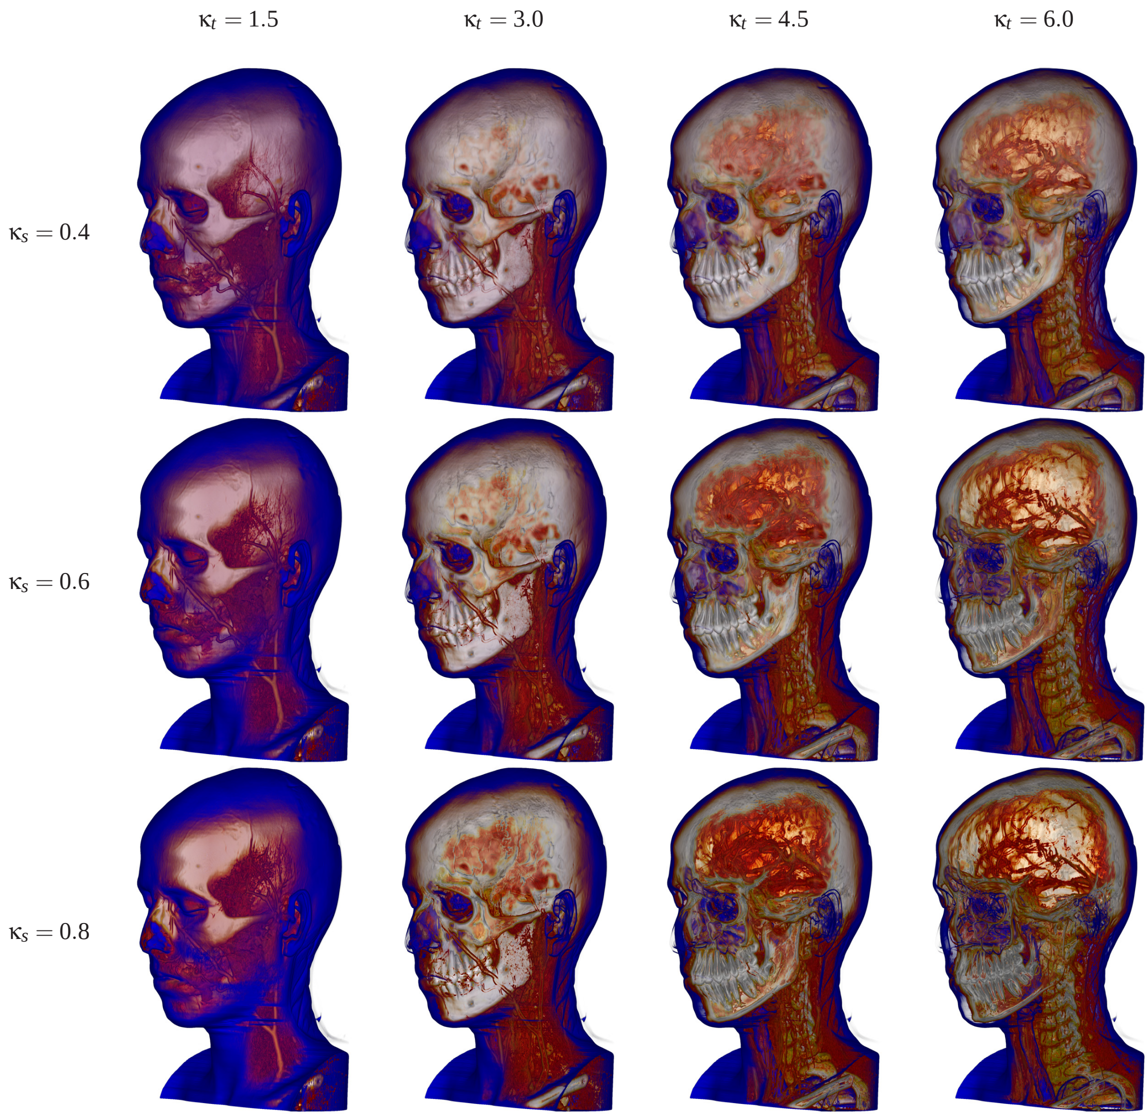

Figure 3: *Context-preserving volume rendering of a contrast-enhanced CT angiography data set using different values for κ_t and κ_s. Columns have the same κ_t value and rows have the same κ_s value.*

These requirements lead us to the following equation for the opacity at each sample position P_i:

$$\alpha(P_i) = \alpha_{tf}(f_{P_i}) \cdot \|g_{P_i}\|^{(\kappa_t \cdot s(P_i) \cdot (1-\|P_i - E\|) \cdot (1-\alpha_{i-1}))^{\kappa_s}} \tag{5}$$

$\|g_{P_i}\|$ is the gradient magnitude and $s(P_i)$ is the shading intensity at the current sample position. A high value of $s(P_i)$ indicates a highlight region and decreases opacity. The term $\|P_i - E\|$ is the distance of the current sample position to the eye point, normalized to the range $[0..1]$, where 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. The influence of the product of these three components is controlled by the two user-specified parameters κ_t and κ_s. The parameter κ_t roughly corresponds - due to the position-dependent term $1 - \|P_i - E\|$ - to the depth of a clipping plane, i.e., higher values reveal more of the interior of the volume. This is the one parameter the user will modify interactively to explore the data set. The effect of modifying κ_s is less pronounced - its purpose is to allow control of the sharpness of the cut. Higher values will result in very sharp cuts, while lower values produce smoother transitions.

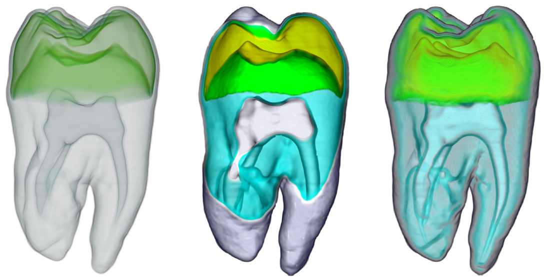

(a) (b) (c)

Figure 4: *CT scan of a tooth rendered with three different techniques.* ***(a)*** *Gradient-magnitude opacity-modulation.* ***(b)*** *Direct volume rendering with clipping plane.* ***(c)*** *Context-preserving volume rendering.*

Figure 3 shows results for different settings of κ_t and κ_s. It can be seen that as κ_t increases more of the interior of the head is revealed - structures on the outside are more and more reduced to the strongest features. An increase in κ_s causes a sharper transition between attenuated and visible regions. A further property of this model is that it is a unifying extension of both direct volume rendering and gradient-magnitude opacity-modulation. If κ_t is set to zero, the opacity remains unmodified and normal direct volume rendering is performed. Likewise, when κ_s is set to zero, the opacity is directly modulated by the gradient magnitude.

4. Results

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 right 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 κ_t and κ_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. Thus, a key advantage of our approach is that it reduces the complexity of transfer function specification. In the following, we present some results achieved with our model in combination with a one-dimensional color transfer function. No segmentation was applied.

(a)

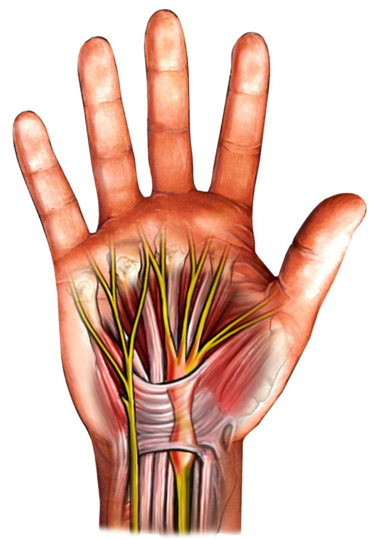

(b)

Figure 5: *Comparing context-preserving volume rendering to illustration.* ***(a)*** *Context-preserving volume rendering of a hand data set.* ***(b)*** *Medical illustration using ghosting.*

Figure 4 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 over-

Figure 6: *Contrast-enhanced CT scan of a leg rendered using context-preserving volume rendering.*

Figure 7: *Contrast-enhanced CT scan of a torso rendered using context-preserving volume rendering.*

lapping transparent structures make the interpretation of the image a difficult task. On the other hand, it is very difficult to place a clipping plane in a way that it does not cut away features of interest. Using context-preserving volume rendering, the clipping depth adapts to features characterized by varying lighting intensity or high gradient magnitude.

Figure 5 and Figure 6 show CT scans of a hand and a leg rendered using our model. These images have a strong resemblance to medical illustrations using the *ghosting* technique, as can be seen by comparing Figure 5 (a) and (b). By preserving certain characteristic features, such as 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.

Finally, Figure 7 shows a CT scan of a human torso. While the image shows many of the features contained in the data set, no depth ambiguities occur as the opacity is selectively varied.

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/vis/adapt/2004_cpvr`

5. Discussion

The context-preserving volume rendering model 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. The two parameters κ_t and κ_s allow intuitive control over the visualization: κ_t is used to interactively browse through the volume, similar to the variation of the depth of a clipping plane; κ_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, 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 κ_t and κ_s, but their values are modulated by a simple function of the scalar value. In Equation 5, κ_t is replaced by $\kappa_t \cdot \lambda_t(f_{P_i})$ and κ_s is replaced by $\kappa_s \cdot \lambda_s(f_{P_i})$. Both λ_t and λ_s are real-valued functions in the range $[0..1]$. For example, the user can specify zero for λ_t to make some regions impenetrable. Likewise, setting λ_s to zero for certain values ensures pure gradient-magnitude opacity-modulation. If one of these functions has a value of one, the corresponding global parameter remains unchanged. Figure 8 (a) shows the visible human male CT data set rendered using just the global parameters, while in Figure 8 (b) bone is made impenetrable by setting λ_t to zero for the corresponding values. Further degrees of freedom of our 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

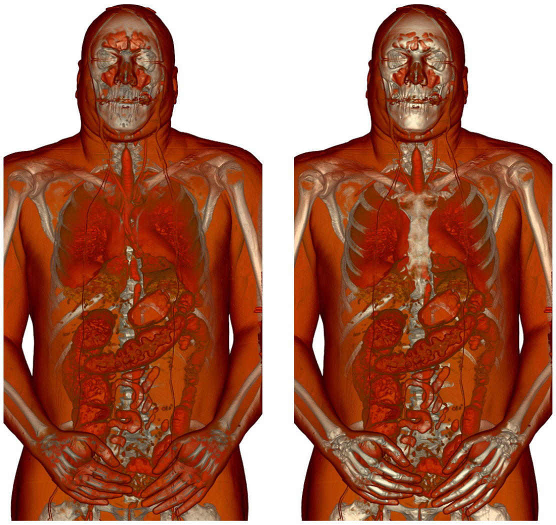

(a) (b)

Figure 8: *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.*

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.

Our method does not pose any restrictions on the type of transfer function used. The model could be applied to modulate the opacity retrieved from a multi-dimensional transfer function without changes. Likewise, the modulation functions λ_t and λ_s could be multi-dimensional. For reasons of simplicity, we have only considered simple one-dimensional transfer functions in this paper.

6. 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 only uses 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 [GBKG04a, GBKG04b]. It could also be used in an implementation using graphics hardware, such as GPU-based ray casting presented by Krüger and Westermann [KW03].

In our implementation, the most costly part of the context-preserving volume rendering model are the exponentiations. However, as with the specular term of the Phong illumination model, it is sufficient to approximate the exponentiation with a function that evokes a similar visual impression. Schlick [Sch94] proposed to use the following function:

$$x^n \approx \frac{x}{n - nx + x} \tag{6}$$

Thus, we can approximate the two exponentiations used by our model in the following way:

$$x^{(a \cdot y)^b} \approx \frac{x \cdot (b + a \cdot y - a \cdot b \cdot y)}{a \cdot y + b \cdot (a \cdot x \cdot y)} \tag{7}$$

This can be efficiently implemented using just 4 multiplications, 2 additions, 1 subtraction, and 1 division. In our implementation, this optimization reduced the cost from a 15% to a 5% increase compared to normal direct volume rendering.

7. Conclusions

The focus of our research was to develop an effective alternative to conventional clipping techniques for volume visualization. It preserves context information when inspecting the interior of an object. The context-preserving volume rendering model is inspired by artistic illustration. 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. Variation of the parameters of the model can then be used to interactively explore the volumetric data set. The achieved results have a strong resemblance to artistic illustrations and, thus, despite their increased information content, are very easy to interpret. Furthermore, as this approach adds little complexity to conventional direct volume rendering, it is well-suited for interactive viewing and exploration. For future work, it might be interesting to alter further properties (e.g., color or compositing mode) in addition to opacity using the presented concept.

We believe that research dealing with the inclusion of low-level and high-level concepts from artistic illustration is very beneficial for the field of volume visualization. The continuation of this research will include the further investigation of the presented approach as well as other illustrative volume rendering techniques in the context of an automated intent-based system for volume illustration [SF91].

8. Acknowledgements

The presented work has been funded by the ADAPT project (FFF-804544). ADAPT is supported by *Tiani Medgraph*, Austria (http://www.tiani.com), and the *Forschungsförderungsfonds für die gewerbliche Wirtschaft*, Austria. See http://www.cg.tuwien.ac.at/research/vis/adapt for further information on this project.

References

[CMH*01] CSÉBFALVI B., MROZ L., HAUSER H., KÖNIG A., GRÖLLER M. E.: Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum 20*, 3 (2001), 452–460.

[ER00] EBERT D. S., RHEINGANS P.: Volume illustration: non-photorealistic rendering of volume models. In *Proceedings of IEEE Visualization 2000* (2000), pp. 195–202.

[GBKG04a] GRIMM S., BRUCKNER S., KANITSAR A., GRÖLLER E.: 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.

[GBKG04b] GRIMM S., BRUCKNER S., KANITSAR A., GRÖLLER M. E.: A refined data addressing and processing scheme to accelerate volume raycasting. *Computers & Graphics 28*, 5 (2004), 719–729.

[HMBG01] HAUSER H., MROZ L., BISCHI G. I., GRÖLLER M. E.: Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 242–252.

[KKH01] KNISS J., KINDLMANN G., HANSEN C.: Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings of IEEE Visualization 2001* (2001), pp. 255–262.

[KVPL04] KONRAD-VERSE O., PREIM B., LITTMANN A.: Virtual resection with a deformable cutting plane. In *Proceedings of Simulation und Visualisierung 2004* (2004), pp. 203–214.

[KW03] KRÜGER J., WESTERMANN R.: Acceleration techniques for GPU-based volume rendering. In *Proceedings of IEEE Visualization 2003* (2003), pp. 287–292.

[KWTM03] KINDLMANN G., WHITAKER R., TASDIZEN T., MÖLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proceedings of IEEE Visualization 2003* (2003), pp. 513–520.

[Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications 8*, 3 (1988), 29–37.

[LM04] LUM E. B., MA K.-L.: Lighting transfer functions using gradient aligned sampling. In *Proceedings of IEEE Visualization 2004* (2004), pp. 289–296.

[LME*02] LU A., MORRIS C. J., EBERT D. S., RHEINGANS P., HANSEN C.: Non-photorealistic volume rendering using stippling techniques. In *Proceedings of IEEE Visualization 2002* (2002), pp. 211–218.

[RE01] RHEINGANS P., EBERT D. S.: Volume illustration: Nonphotorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 253–264.

[Sch94] SCHLICK C.: A fast alternative to Phong's specular model. In *Graphics gems IV*, Heckbert P., (Ed.). Academic Press, 1994, pp. 385–387.

[SF91] SELIGMANN D. D., FEINER S. K.: Automated generation of intent-based 3d illustrations. In *Proceedings of ACM Siggraph 1991* (1991), pp. 123–132.

[VKG04] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven volume rendering. In *Proceedings of IEEE Visualization 2004* (2004), pp. 139–145.

[WEE02] WEISKOPF D., ENGEL K., ERTL T.: Volume clipping via per-fragment operations in texture-based volume visualization. In *Procceedings of IEEE Visualization 2002* (2002), pp. 93–100.

[WEE03] WEISKOPF D., ENGEL K., ERTL T.: Interactive clipping techniques for texture-based volume visualization and volume shading. *IEEE Transactions on Visualization and Computer Graphics 9*, 3 (2003), 298–312.

[WK95] WANG S. W., KAUFMAN A. E.: Volume sculpting. In *Proceedings of the Symposium on Interactive 3D Graphics 1995* (1995), pp. 151–156.

[ZDT04] ZHOU J., DÖRING A., TÖNNIES K. D.: Distance based enhancement for focal region based volume rendering. In *Proceedings of Bildverarbeitung für die Medizin 2004* (2004), pp. 199–203.