

Importance-Driven Volume Rendering

Ivan Viola

Armin Kanitsar

Meister Eduard Gröller*

Institute of Computer Graphics and Algorithms
Vienna University of Technology, Austria

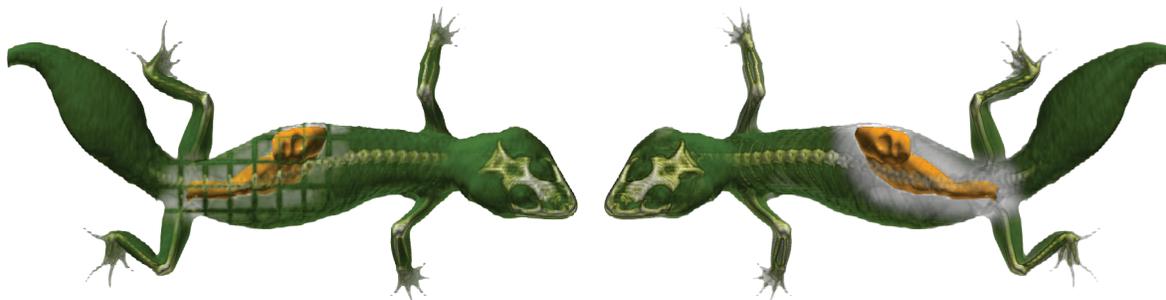


Figure 1: Leopard Gecko dataset: Importance-driven volume rendering of interior structures.

ABSTRACT

This paper introduces importance-driven volume rendering as a novel technique for automatic focus and context display of volumetric data. Our technique is a generalization of cut-away views, which – depending on the viewpoint – remove or suppress less important parts of a scene to reveal more important underlying information. We automatize and apply this idea to volumetric data.

Each part of the volumetric data is assigned an *object importance* which encodes visibility priority. This property determines which structures should be readily discernible and which structures are less important. In those image regions, where an object occludes more important structures it is displayed more sparsely than in those areas where no occlusion occurs. Thus the objects of interest are clearly visible. For each object several representations, i.e., *levels of sparseness*, are specified. The display of an individual object may incorporate different levels of sparseness. The goal is to emphasize important structures and to maximize the information content in the final image.

This paper also discusses several possible schemes for *level of sparseness* specification and different ways how *object importance* can be composited to determine the final appearance of a particular object.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing algorithms;

Keywords: view-dependent visualization, volume rendering, focus+context techniques, level-of-detail techniques, non-photorealistic techniques

*{viola|kanitsar|meister}@cg.tuwien.ac.at,
<http://www.cg.tuwien.ac.at/home/>

1 INTRODUCTION

The importance of volumetric visualization in medical applications has been increasing significantly over the last years. In the future three dimensional visualization will be an essential tool for medical diagnosis. Due to the rapid development of high precision imaging modalities the amount of data keeps increasing. Therefore the need for emphasizing important structures raises.

Various medical applications exist where the size of the structures of interest is comparably small in relation to the overall data. Examples are examinations of tumors in the kidneys, lesions inside the liver and findings of lung nodules. In addition to that, performing diagnostic examinations is a complex task. Different features and properties of the anatomical tissues have to be taken into account. Especially not only the size and the shape of pathologies are of interest, but also their spatial position and vicinity to other anatomical structures. Hence, from a computer science point of view we are dealing with a focus and context task.

The investigation of liver lesions illustrates the medical requirements on the applied visualization method. Radiologists need to see the tumor from several directions in order to estimate the shape of the lesion. Computer aided diagnosis tools provide precise information about the size calculated from volumetric segmentation. The exact spatial position of arteries in close vicinity is very important in order to determine which liver segments to remove in a possible subsequent surgical treatment. Therefore it is necessary to visualize three different structures: the tumor, the vessel tree of the liver, and – of course – the liver parenchyma.

Visualizing these structures simultaneously – probably also including parts of the surrounding anatomy for orientation purposes – results in objects occluding each other. Known methods classify objects within the dataset independently from the viewpoint. This limits viewpoint positions and viewing angles to a range, where the important structures are not occluded by other objects. Using clipping geometries to reveal important interior structures eliminates less important objects also in those viewing situations where it would not be necessary. The application of different optical prop-

erties and rendering techniques (i.e., silhouette rendering) eases the problem only to a certain degree. From certain viewpoints several objects of less importance may overlap and largely obscure important features. Furthermore the fine-tuning of rendering parameters is a time consuming process not suitable for rapid clinical use.

These and similar visualization tasks can be resolved by *importance-driven volume rendering (IDVR)*. The tumor and the vascular tree in close vicinity are the most important objects, the liver tissue and the surrounding anatomy (bones, aorta, skin) are of lower importance but still helpful for orientation purposes. Similar to anatomical drawings with IDVR the interesting structures are clearly visible from varying positions and different viewing angles. Occluding objects are rendered more sparsely or suppressed totally.

The main contribution of this paper is importance-driven volume rendering as a model for automatic focus and context rendering. The proposed method overcomes the problem of occlusion within the volume, which happens when using any kind of view-independent classification. As opposed to previous approaches the optical properties of the proposed technique are not constant for an entire object. Depending on the viewing situation the *level of sparseness* varies dynamically. In order to visually emphasize features with the highest importance, objects between these features and the viewpoint are rendered very sparsely. Interesting objects are represented more densely to see most of the detail. Irrespective of the viewpoint there is always a clear view on the most important structures. If no occlusion occurs, even the less important objects can be rendered densely. This enables an automatic generation of images with maximal visual information.

In Figure 2 an anatomical illustration of the human abdomen [12] and a result of our technique is presented. In this case the internal structures are classified with a high importance value so that structures between the viewpoint and the important values are simply cut away automatically.

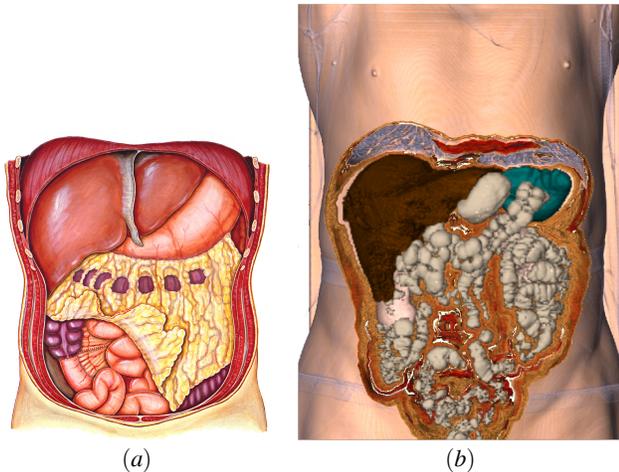


Figure 2: Comparison between artistic medical illustration of the abdomen (a) and our method (b). Illustration image is courtesy to Howell MediGraphics.

The paper is organized as follows: Section 2 describes previous work related to importance-driven volume rendering. Section 3 explains the basic idea of the proposed model. Section 4 discusses the principal components of the model and depicts their impact on the resulting visualization. In Section 5 we describe the test datasets, implementation environment and rendering results. Finally we draw conclusions, summarize the paper in Section 6 and propose future work in this field in Section 7.

2 RELATED WORK

Scientific work related to our model can be divided into several categories. We will first discuss methods that use advanced transfer function specification for feature enhancement and use semi-automatic transfer function specification. These classification methods are viewpoint-independent. Our work enables automatic focus+context visualization, where the information about the viewpoint is taken into account. We therefore point out some previous focus+context approaches. Afterwards we review various rendering techniques that are interesting for the *level of sparseness* specification. Finally we briefly mention other graphics research fields, where importance information plays a significant role.

Feature Classification: The most typical feature classification in volume visualization is the *transfer function* specification that maps sample density values to optical properties, i.e., color and opacity. Transfer functions with density as single input parameter are also denoted as one-dimensional transfer functions. There has also been a lot of research on multi-dimensional transfer functions. The basic idea is to incorporate first and second derivatives of the density into the transfer function design [11, 17]. It is possible to assign optical properties based on gradient and curvature values, so for example object boundaries are classified differently from homogeneous regions. Taking into account first and second derivatives enables the semi-automatic generation of transfer functions [16]. An interesting approach was presented by Hauser and Mlejnek [9] for multidimensional 3D flow data. They use the *Degree-of-Interest* function that maps the user interest to optical properties.

Focus+Context Rendering: A lot of work has been done in the field of focus+context visualization. Viewpoint-dependent distortion of three-dimensional data [2] for example highlights data by dedicating more space to it. Distortions are applied to abstract graphs in order to clearly see interesting graph nodes. An interesting idea is also to include the *distance to focal point* property into the volume rendering pipeline [29]. This allows to change the optical properties according to the distance to the focal point. Using this technique several expressive focus+context effects can be achieved. Gaze-directed volume rendering takes the observer’s viewing direction into account to increase the rendering performance [19]. The volume dataset is rendered in different resolutions. According to the viewing direction the focal region is represented in full resolution, while the other parts are rendered in lower resolution.

Sparse Representation: The graphics community has been inspired by artists to reduce a visual representation just to show features of interest. This is in contrast to traditional photorealistic approaches. Non-photorealistic methods show features in a sparse way exploiting human imagination. The display of contours is a popular method to represent context information in volume visualization [3, 23]. Outlines are often sufficient to roughly understand the shape and can be combined with other rendering techniques like direct volume rendering or maximum intensity projection [10]. To make a contour representation more dense *suggestive contours* can be introduced [5]. Additionally to real contours also contours from virtual viewpoints close to the current view are rendered. Also *pen-and-ink* techniques convey good shape information. Pen-and-ink styles in combination with traditional volume rendering have already been applied for focus+context rendering in volume visualization [27]. This is up to a certain degree similar to combining curvature-directed strokes with iso-surface rendering [13]. This approach was proposed for rendering structures that are completely enclosed by other objects. The interior structures are rendered fully opaque, while the enclosing objects are represented by a set of curvature-directed lines. The visibility of interior structures can also be modified by dynamic changes in transparency of the outer shape. Dynamic transparency is also used in user interface design [8].

Cut-Away Views: Cut-away illustrations are another way to represent nested objects. The popularity of this technique is demonstrated by the fact that it can be found in almost all books with technical or medical illustrations. In volume visualization this technique is also known as *volume cutting* [22]. Diepstraten et al. [6] describe how various cut-away techniques can be achieved automatically. Streamarrows were presented by Löffelmann et al. [20] for visualizing complex dynamical systems. They use arrows as a basic element for cutting the stream surface. This allows to see through the surface and perceive other surfaces or structures behind.

Importance-Driven Rendering: Importance as a rendering parameter has already been investigated in other graphics fields. In global illumination the term *importance* is used quite often. Importance sampling for example is a popular method to accelerate convergence of Monte-Carlo integration by faster reducing the variance. The basic idea distributes sample points non-uniformly [26, 14]. This idea of has been recently researched also for volume rendering [7, 4]. For radiosity techniques importance can also encode the potential visibility of surfaces [24] or the prominence of a patch as a light source. Importance is also proposed as an additional user-defined parameter for halftoning [25]. All these methods employ *importance* in different contexts and meanings. But all of them use importance as an additional dimension to improve the behavior of traditional approaches. In the following we discuss how *importance* is applicable also for volume rendering.

3 IMPORTANCE-DRIVEN VOLUME RENDERING

In volume visualization we are often dealing with the problem that interesting structures are partly or completely occluded by surrounding tissue. This is hard to resolve by traditional view-independent approaches, such as transfer function specification. We propose a viewpoint-dependent model that removes unwanted occlusions automatically and maximizes the information content of the final image.

Traditionally objects within the volume dataset are classified by optical properties like color and opacity. We additionally assign objects another dimension, which describes their *importance*. Importance encodes which objects are the most interesting ones and have the highest priority to be clearly visible. Each object is therefore weighted by a positive scalar value called *object importance*. During the rendering stage, the model evaluates the visibility of each object according to its importance. If less important objects are occluding features that are more interesting, the less important ones are rendered more sparsely, e.g., more transparently. If the same object does not cause any unwanted occlusions in other regions of the image, it is rendered more densely, e.g., opaque, in order to see its features more clearly. This allows to see all interesting structures no matter if they are occluded or not, and the less important parts are still visible as much as possible. Instead of using constant optical characteristics, which are independent from the viewpoint, we use several *levels of sparseness* for each object. We do not assign a single optical characteristic, but several characteristics with smooth transitions inbetween. These multiple levels of sparseness allow the object to continuously change its visual appearance from a very dense representation to a very sparse one. Which level of sparseness will be chosen is dependent on the importance of the particular object and the importance of hidden objects behind. The level of sparseness thus may continuously vary within a single object. Also depending on the viewpoint the same part of an object may be represented with different levels of sparseness. To determine the sparseness level for each object or parts thereof the rendering pipeline requires an additional step, which we call *importance compositing*. This step evaluates the occlusion, takes the importance factor of each object into account and assigns to each object particular *levels of sparseness*. The final synthesis results in images

with maximal visual information with respect to the predefined object importance.

For the sake of clarity the controlling components of importance-driven volume rendering are now briefly summarized:

Object importance describes the visibility priority of each object within the volume dataset. It is a positive scalar value, which is constant for the whole object.

Levels of sparseness are various different representations of a particular object from the most dense to the most sparse one. Sparseness is defined in terms of how much the display of an object takes up screen estate. For an iso-surface a point-cloud or a wire-frame display is a very sparse representation. The display of filled opaque polygons is a very dense representation of the same iso-surface. In case of volumetric data each sample within the volume is classified to have a color and opacity. In direct volume rendering (DVR) these values are composited together to synthesize the final image. The sparseness can in this case modulate the opacity of object samples so that the most interesting samples are opaque while other samples are more transparent. This is a simple and straightforward example of levels of sparseness. More elaborate schemes are discussed in Section 4.2.

Importance Compositing is an additional rendering step that assigns to each part of an object a specific *level of sparseness*. A very simple importance compositing scheme is maximum importance projection (MImP). For each image area only the object with highest importance is visible. The remaining objects are shown with the highest level of sparseness, i.e., fully transparent. This results in cut-out views of important structures. As another example, we can perform importance compositing in a similar way as compositing of color and opacity in the traditional DVR approach. Instead of compositing optical properties, the *object importance* values are accumulated. We will discuss different compositing modes in Section 4.1.

The interrelationship between the above mentioned components is depicted in Figure 4. The *importance compositing* is done similar to the DVR approach. For each ray the compositing evaluates object occlusions and assigns the corresponding *level of sparseness* to each object. Object importance is presented in the sense that it is mapped to object visibility in the result image. This causes different rendering settings for the context object 0 in the area of the image which is covered by the focus object 1. The difference between traditional volume rendering and importance-driven volume rendering is clearly visible in Figure 3. In the traditional approach it is necessary to reduce the opacity of occluding objects globally. Importance-driven rendering assigns a higher sparseness factor only to the area where occlusion occurs.



Figure 3: Comparison between traditional volume rendering (top) and importance-driven volume rendering (bottom).

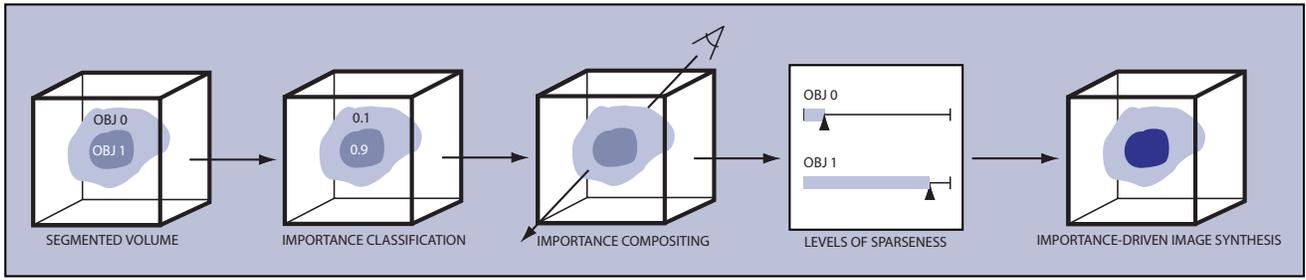


Figure 4: Stages of the importance-driven volume rendering pipeline.

4 IMPORTANCE AND SPARSENESS FACTORS

In the previous section we have explained the basic idea of *importance-driven volume rendering* and its components. For the moment we assume that object importance is just a scalar value. Therefore in this chapter we concentrate on the two remaining components, i.e., importance compositing and levels of sparseness.

4.1 Importance Compositing

Importance compositing is an additional pass added to the traditional volume-rendering pipeline. It determines the level of sparseness for each object or a part thereof in order to preserve important features. There are many possibilities conceivable how to perform importance compositing. In the following we will discuss two simple methods of importance compositing which are derived from compositing optical properties through ray casting of volume data.

Maximum Importance Projection

Maximum Intensity Projection (MIP) [21] is a simple and fast volume rendering approach. It is applicable for sparse data where important information has high intensity values like contrast-media enhanced blood vessels. With MIP compositing reduces to selecting the highest intensity value along a ray. Intensities are encoded as gray values to produce the final image.

Analogous to Maximum Intensity Projection we propose *Maximum Importance Projection (MImP)*. For each ray the object with highest importance along the ray is determined. This object is displayed densely. All the remaining objects along the ray are displayed with the highest level of sparseness, i.e., fully transparent. With *MImP* structures are either rendered using the most dense representation or they are not rendered at all.

With Maximum Intensity Projection the spatial arrangement of structures is not readily apparent. MImP has a similar problem which we alleviate as follows. The image area, where the object is projected onto is denoted as object *footprint*. With MImP the footprint is exactly the image region where only the object is visible. One can consider MImP as a cut-away view, where the space in front of the most important object is simply clipped. The clipping object is a translational sweep with the footprint as cross section (general cylinder). One can now modify this cylinder to obtain a clipping frustum. This is achieved by scaling up the footprint during the translational sweep towards the viewer. This produces a countersink clipping geometry. Figure 5 illustrates the difference between the above described cylindrical and a conical MImP in 2D. The conical MImP is easily realized during rendering ray traversal by changing the starting point during ray traversal for those rays intersecting the side faces of the clipping frustum. Figure 6 shows images to compare both approaches.

The countersink geometry, respectively the ray starting points are computed from the footprint of the focus object. The footprint contains depth information of the focus objects' *last-hit* sample for each ray along the viewing direction. This information is used for performing the cut-out. For cylindrical MImP we simply skip ray samples that belong to the context object, until the focus objects' last hit depth is reached. For the conical MImP we need to enlarge the footprint to build the conical shape. This is possible to realize using image processing operators on the depth image, where the intensity encodes the depth of the entry point. The depth-footprint is processed by a 2D chamfer distance transform [1]. The highest depth value of the footprint is calculated. The depth buffer of ray starting points that contribute to the countersink e_i are calculated from the highest depth value e_{max} , the slope s_c of the countersink and the distance value d_i at pixel i as shown in Equation 1.

$$e_i = e_{max} - d_i * s_c \quad (1)$$

To correctly simulate the cut-out it is necessary to change the gradient vector at the entry points that contribute to the countersink effect. Two components of the gradient are estimated from the gradient information of the 2D distance field. The z component is constant, i.e., the *slope* of the countersink frustum.

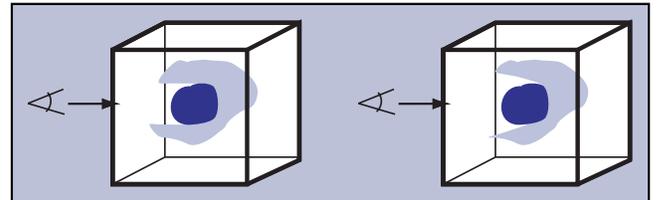


Figure 5: Maximum Importance Projection. Illustration of cylindrical (left) and conical countersink (right).

Average Importance Compositing

The second approach of importance compositing takes into account all the objects along a ray. The influence of an individual object is hereby independent from the number of ray samples within the object. Object o has an importance value I_o . Ray r is intersected by n_r objects. The level of sparseness S_o of a particular object o at ray r is equal to the fraction of its own importance and the sum of the importance of all the intersected objects:

$$S_o = \frac{I_o}{\sum_{i=1}^{n_r} I_i} \quad (2)$$

Average importance compositing does not completely remove the less important objects as with MImP. The sparseness factors are

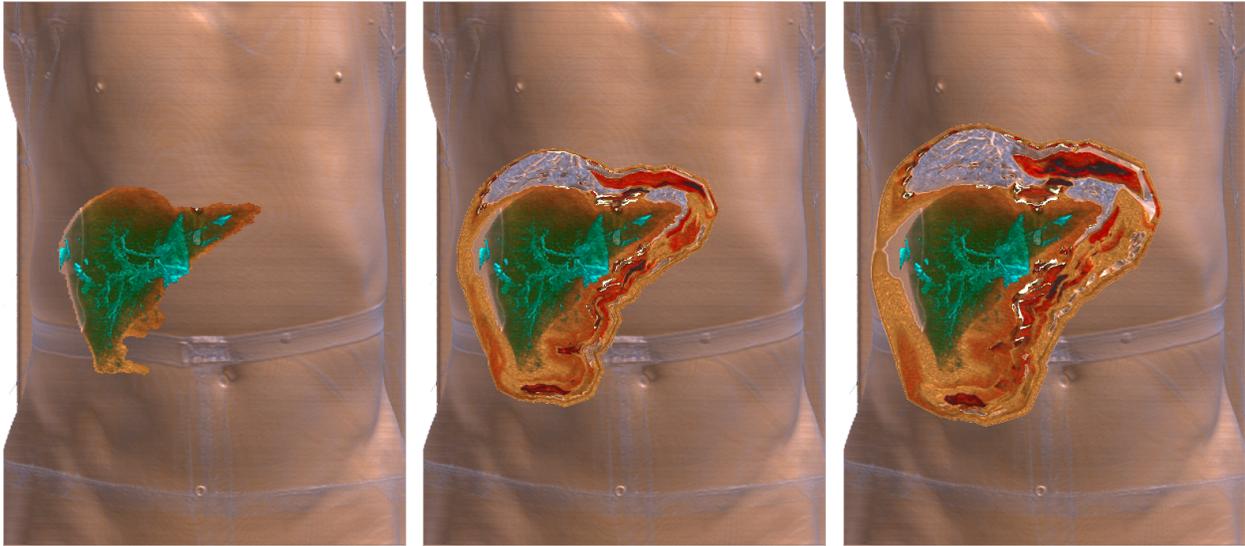


Figure 6: Maximum Importance Projection. Cylindrical MlMP (left) does not clearly show the spatial relationships between focus and context objects, i.e., the focus object appears in front of the context object. Conical MlMP corrects this visual artifact. Different slope factors define the *countersink* geometry (center and right).

estimated according to the given importance. This allows a very sparse representation of the occluding object maybe just to see its high curvature areas and to clearly see the important object behind it. The compositing can be evaluated using the binary footprint image of each object. Then for each image region we get the sparseness factors of the corresponding objects. For each pixel of the final image we perform a lookup to each object footprint. Object importance values of the covering objects are summed together. The sparseness factor is estimated through division of the object importance by the evaluated sum. Unfortunately this approach generates sharp transitions with visible *jaggy* object boundaries. To achieve smooth transitions, we let each footprint *grow* and smooth its boundaries. The footprint image is in this case grayscale instead of binary. The final importance compositing is analogue to the above mentioned approach according to Equation 2. The object importance is additionally always multiplied by the footprint value in the range of $[0, 1]$. Footprint values below one are part of the *transition* area. Figure 7 gives an example of average importance compositing.

4.2 Levels of Sparseness

The importance ray traversal determines for each object respectively parts thereof the visibility on the rendered image by determining the level of sparseness. This section describes three simple schemes how to achieve smooth transitions of object visibility.

These schemes are compared in Figure 8, which shows the *Leopard Gecko* dataset, where the interior part was pre-segmented. The series of images illustrates how the context area in front of the focus object smoothly varies from a dense to a sparse representation.

Color and Opacity Modulation

A direct control of optical properties is the first approach to modify the visual prominence of a particular object. With increasing sparseness the object becomes more transparent in order to show the more important underlying data. This approach is widely used in transfer function specification.

Interesting results can be achieved by controlling color saturation with the level of sparseness factor. Color is a very important

visual cue for visualization in general. Highly saturated colors are attracting the observer's focus more than colors close to gray. The level of sparseness can therefore be expressed also in the saturation of the color. Changing only the saturation, however, does not increase the visibility of occluded objects. Therefore it is necessary to change the color and opacity values at the same time. Different visual appearances within the same object can cause misinterpretations. Therefore smooth transitions between different levels of sparseness have to be applied. A smooth modulation of the optical properties is shown in Figure 8 (top).

Screen-Door Transparency

Screen-door transparency is a well-known strategy to simulate transparency. The semi-transparent view of an object behind a screen-door is simulated as follows: A screen-door consists of a mesh and holes inbetween. The mesh fully occludes the object behind, whereas it is fully or partly visible through the holes. From a certain distance holes and mesh blend together to produce a semi-transparent view. We use an analogue idea to define levels of sparseness. The volumetric dataset consists of voxels. The level of sparseness determines which voxels should be rendered and which not. The distribution of visible voxels is uniform and is forming a wireframe-like structure. The impact of increasing sparseness is shown in Figure 8 (middle).

Volume Thinning

Volume Thinning proceeds as follows: Voxels of an object are sorted according to two sorting keys. The first sorting key is gradient magnitude, the second sorting key is curvature magnitude of the iso-surface through the voxel. Reducing the sparseness factor according to gradient magnitude has the effect that the volume is continuously reduced to fewer and fewer strong iso-surfaces. As soon as only few iso-surfaces remain the reduction proceeds according to curvature magnitude. This has the effect that the iso-surfaces gradually dissolve and in the end (most sparse representation) only few high curvature areas remain. Figure 8 (bottom) illustrates visibility reduction through volume thinning.



Figure 7: Average importance compositing: The average compositing is shown under different viewpoint settings in combination with modulating optical properties.

5 RESULTS

We show results of our method on two datasets. The *Leopard Gecko* dataset is of resolution $512 \times 512 \times 87$. The *Monster Study* dataset has been downsampled to the half of its full resolution, i.e., $256 \times 256 \times 610$. Both datasets are using pre-segmented objects. The *Monster Study* has all important organs pre-segmented. Therefore it is optimal for automatic generation of medical illustrations.

The performance of the current implementation is not interactive. The goal of the implementation was to verify the proof of concept rather than performance optimizations. The model was integrated into the J-Vision [15] medical workstation as a plugin.

The rendering results of the *Monster Study* dataset (Figures: 2 and 6) show the applicability of Maximum Importance Projection in medical imaging. The images of the *Leopard Gecko* illustrate the effect of average importance compositing using different levels of sparseness (Figures: 1, 3, 7, and 8).

To fully appreciate the strengths of importance-driven volume rendering viewpoint changes are essential. This is best illustrated with animation sequences, which are available at http://www.cg.tuwien.ac.at/research/vis/adapt/2004_idvr/.

6 SUMMARY AND CONCLUSIONS

In this paper we have proposed importance-driven volume rendering as a view-dependent model for automatic focus+context volume visualization. A new factor to the traditional volume rendering pipeline is introduced, i.e., the importance dimension. According to the importance and viewpoint settings each object is rendered in order to maximize the visual information. This method allows to see structures within the volume as dense as possible. A sparse representation is chosen only if other, more important structures are occluded.

We have discussed three schemes for levels of sparseness. This factor controls the optical properties or the amount of visible elements of the volume. Smooth opacity changes work well in combination with desaturation. The amount of visible volume elements can be distributed uniformly over the whole volume, or the first- and second-order derivatives can be used for visibility distribution.

7 FUTURE WORK

The paper opens a lot of opportunities for possible research areas. We assume that the objects within the volume data are pre-segmented and the importance is assigned manually by the user. An open issue is, how to do the feature selection and importance

assignment automatically. Various automatic feature detection approaches can be integrated into the model to select the important features without additional user interaction. Another idea can be to assign the importance values according to closeness to one or several focal points, similarly to distance-based enhancement for focal region [29]. The importance dimension controls in this case the transition from one level of sparseness to another.

The paper has presented various levels of sparseness schemes. The continuous transition from dense to sparse representations for volumetric data is a wide area of research. In polygonal rendering levels of sparseness are widely used. The most sparse representation is a set of points, another representation is a wireframe display, and the most dense display is a surface representation. Volume graphics does not yet have such a variety, which shows the need for research in this area.

The third factor of importance-driven volume rendering is importance compositing. The paper presents simple compositing schemes derived from ray-casting approaches. The next step are compositing schemes that incorporate first- and second order derivatives to preserve object boundaries. These parts are then considered as more important and a dense representation is chosen.

The conical *MImP* is a compositing scheme that uses the cut-out illustration technique to improve perception of the spatial relationships. More elaborate approaches for intelligent automatic cut-out generations are to be researched more intensively. In cut-away views sometimes the borderline of the cutted regions is emphasized (e.g., through thick lines, or zig-zag lines). Automatically emphasizing these transition zones is also an open problem.

Each viewpoint brings out only a fraction of the entire information encoded in the data set. How to estimate viewpoint entropy and how to automatically determine optimal viewpoints [28] is another, not yet researched, area for volume data.

8 ACKNOWLEDGMENTS

The work presented in this publication has been funded by the ADAPT project (FFF-804544). ADAPT is supported by *Tiani Medgraph*, Vienna (<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.

The authors would like to thank Stefan Bruckner and Sören Grimm for providing the software ray-caster. The *Monster Study* dataset is courtesy of Tiani Medgraph and the *Leopard Gecko* is courtesy of University of Veterinary Medicine Vienna. The medical illustration of the human abdomen is courtesy of Howell Medi-Graphics [12].

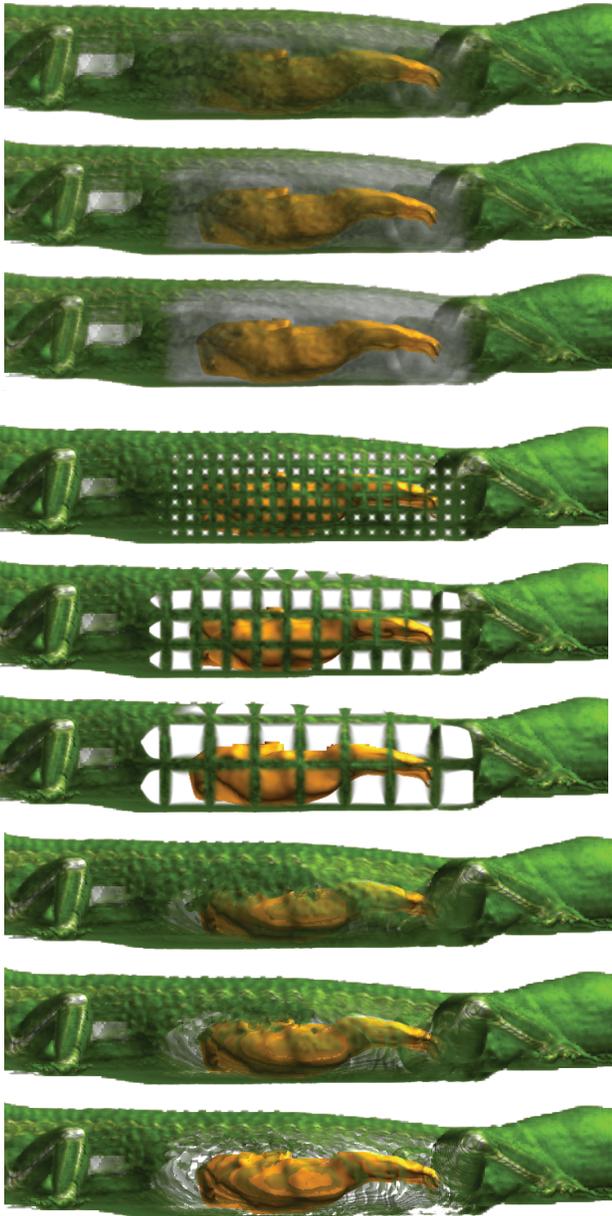


Figure 8: Changing levels of sparseness. Top three rows: opacity modulation and color saturation modulation. Middle three rows: screen-door transparency. Bottom three rows: volume thinning. Images display levels of sparseness with factors 0.75, 0.5, and 0.25.

REFERENCES

[1] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344–371, 1986.

[2] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. Distortion viewing techniques for 3-dimensional data. In *Proceedings of IEEE Symposium on Information Visualization '96*, pages 46–53, 1996.

[3] B. Csébfalvi, L. Mroz, H. Hauser, A. König, and M. E. Gröller. Fast visualization of object contours by non-photorealistic volume rendering. In *Proceedings of EUROGRAPHICS '01*, pages 452–460, 2001.

[4] B. Csébfalvi and L. Szirmay-Kalos. Monte carlo volume rendering. In *Proceedings of IEEE Visualization '03*, pages 449–456, 2003.

[5] D. DeCarlo, A. Finkelstein, and S. Rusinkiewicz A. Santella. Suggestive contours for conveying shape. In *Proceedings of ACM SIGGRAPH '03*, pages 848–855, 2003.

[6] J. Diepstraten, D. Weiskopf, and T. Ertl. Interactive cutaway illustrations. In *Proceedings of EUROGRAPHICS '03*, pages 523–532, 2003.

[7] S. Grimm, S. Bruckner, A. Kanitsar, and E. Gröller. VOLUME doTS as a point-based representation of volumetric data. In *To appear in Proceedings of EUROGRAPHICS '04*, 2004.

[8] C. Gutwin, J. Dyck, and Ch. Fedak. The effects of dynamic transparency on targeting performance. In *Proceedings of Graphics Interface '03*, pages 105–112, 2003.

[9] H. Hauser and M. Mlejnek. Interactive volume visualization of complex flow semantics. In *Proceedings of VMV '03*, pages 191–198, 2003.

[10] H. Hauser, L. Mroz, G. Bisch, and M. E. Gröller. Two-level volume rendering-fusing MIP and DVR. In *Proceedings of IEEE Visualization '00*, pages 211–218, 2000.

[11] J. Hladůvka, A. König, and E. Gröller. Curvature-based transfer functions for direct volume rendering. In *Proceedings of SCCG '00*, pages 58–65, 2000.

[12] Howell MediGraphics web site, <http://www.medgraphics.com/>, 2004.

[13] V. Interrante, H. Fuchs, and S. Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *Proceedings of IEEE Visualization '96*, pages 211–218, 1996.

[14] H. Wann Jensen. Importance driven path tracing using the photon map. In *Proceedings of Rendering Techniques '95*, pages 326–335, 1995.

[15] J-Vision medical workstation, Tiani MedGraph web site, <http://www.tiani.com/>, 2004.

[16] G. Kindlmann and J. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of IEEE Symposium on Volume Visualization '98*, pages 79–86, 1998.

[17] J. Kniss, G. Kindlmann, and Ch. Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings of IEEE Visualization '01*, pages 255–262, 2001.

[18] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8:29–37, 1987.

[19] M. Levoy and R. Whitaker. Gaze-directed volume rendering. In *Proceedings of Symposium on Interactive 3D Graphics '90*, pages 217–223, 1990.

[20] H. Löffelmann, L. Mroz, and E. Gröller. Hierarchical streamarrows for the visualization of dynamical systems. In *Eurographics Workshop on Visualization in Scientific Computing '97*, pages 155–164, 1997.

[21] L. Mroz, H. Hauser, and M. E. Gröller. Interactive high-quality maximum intensity projection. In *Proceedings of EUROGRAPHICS '00*, pages 341–350, 2000.

[22] B. Pflesser, U. Tiede, and K. H. Höhne. Towards realistic visualization for surgery rehearsal. In *Computer Vision, Virtual Reality and Robotics in Medicine, Proc. CVRMed '95*, pages 487–491, 1995.

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

[24] B. Smits, J. Arvo, and D. Salesin. An importance-driven radiosity algorithm. In *Proceedings of ACM SIGGRAPH '92*, pages 273–282, 1992.

[25] L. Streit and J. Buchanan. Importance driven halftoning. In *Proceedings of EUROGRAPHICS '98*, pages 207–217, 1998.

[26] L. Szirmay-Kalos, B. Csébfalvi, and Werner Purgathofer. Importance driven quasi-random walk solution of the rendering equation. In *Proceedings of WSCG '98*, pages 379–385, 1998.

[27] S. Treavett and M. Chen. Pen-and-ink rendering in volume visualisation. In *Proceedings of IEEE Visualization '00*, pages 203–210, 2000.

[28] P. P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Viewpoint selection using viewpoint entropy. In *Proceedings of VMV '01*, pages 273–280, 2001.

[29] J. Zhou, A. Döring, and K. Tönnies. Distance based enhancement for focal region based volume rendering. In *Proceedings of Bildverarbeitung für die Medizin '04*, pages 199–203, 2004.