

Importance-Driven Feature Enhancement in Volume Visualization

Ivan Viola, Armin Kanitsar, and M. Eduard Gröller, *Member, IEEE*

(Invited Paper)

Abstract—This paper presents importance-driven feature enhancement as a technique for the automatic generation of cut-away and ghosted views out of volumetric data. The presented focus+context approach removes or suppresses less important parts of a scene to reveal more important underlying information. However, less important parts are fully visible in those regions, where important visual information is not lost, i.e., more relevant features are not occluded.

Features within the volumetric data are first classified according to a new dimension denoted as *object importance*. This property determines which structures should be readily discernible and which structures are less important. Next, for each feature various representations (*levels of sparseness*) from a dense to a sparse depiction are defined. Levels of sparseness define a spectrum of optical properties or rendering styles. The resulting image is generated by ray-casting and combining the intersected features proportional to their importance (*importance compositing*).

The paper includes an extended discussion on several possible schemes for *levels of sparseness* specification. Furthermore different approaches to *importance compositing* are treated.

Index Terms—view-dependent visualization, volume rendering, focus+context techniques, level-of-detail techniques, illustrative techniques

I. INTRODUCTION

THE relevance of volume visualization in medical applications has been increasing over the last years. Three-dimensional visualization is becoming an essential tool for medical diagnosis and operation planning. Due to the rapid development of high-precision imaging modalities the amount of data is steadily increasing. The amount of relevant information is often relatively small as compared to the overall amount of acquired data. Therefore these small, interesting features have to be visually emphasized. Examples are tumors in the kidneys, lesions inside the liver, and lung nodules. Diagnostic examinations are complex tasks, where properties

of the anatomical tissues have to be taken into account. In addition to the size and shape of pathologies also their spatial position and vicinity to other anatomical structures is of interest. Hence, from a computer science point of view it is a focus+context task.

The detection of liver lesions illustrates the medical requirements on the applied visualization method. Medical experts need to see the tumor from several directions in order to estimate the shape of the lesion. Furthermore the spatial position of arteries in close vicinity is very important in order to determine which liver segments must be removed in a possible subsequent surgical treatment. The visualization task is to display three different structures: the tumor, the vessel tree of the liver, and the liver parenchyma. However, displaying these structures simultaneously results in objects occluding each other. Traditional techniques classify objects within the dataset independently from the viewpoint. The global setting limits viewpoint positions and viewing angles to a range, where the important structures are not occluded by other objects. One possibility is to use clipping planes. Such an approach eliminates less important objects also in those viewing situations, where it would not be necessary. Different optical properties and rendering techniques (i.e., silhouette rendering) ease the problem only to a certain degree and these settings are applied globally. Beside this the fine-tuning of rendering parameters is a time consuming process not suitable for rapid clinical use.

Medical tasks such as visualizing liver lesions can be resolved by *importance-driven volume rendering* (IDVR) [28]. The tumor and the vascular tree in close vicinity are the most important features, the liver tissue and the surrounding anatomy (bones, aorta, skin) are of lower importance but still helpful for orientation purposes. With IDVR the interesting structures are clearly visible from different viewing angles. Occluding objects are rendered more sparsely or suppressed entirely.

The main contribution of this paper is importance-driven feature enhancement as an approach to automatic focus+context volume rendering. The proposed method overcomes the problem of occlusions within the volume,

The authors are with the Institute of Computer Graphics and Algorithms at the Vienna University of Technology, Austria. E-mail: {viola|kanitsar|meister}@cg.tuwien.ac.at

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 estimated *level of sparseness* varies dynamically. In order to visually emphasize features with the highest importance, occluding objects between these features and the viewpoint are rendered sparsely. Interesting objects are represented more densely to see most of the details. If no occlusion occurs, even the less important features can be rendered densely. This enables an automatic generation of images with maximal visual information.

In Figure 1 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 features are simply cut away automatically.

The paper is organized as follows: Section II describes previous work related to importance-driven volume rendering. Section III explains the basic idea of the proposed model. Sections IV and V discuss the principal components of the model, i.e., *importance compositing* and *levels of sparseness*, and depict their impact on the resulting visualizations. In Section VI we describe the test datasets and show various rendering results applicable for medical and flow visualization. Finally we draw conclusions, summarize the paper in Section VII and propose future work in Section VIII.

II. RELATED WORK

Scientific work related to our model can be divided into several categories. First, methods that use advanced and semi-automatic transfer function specification for feature enhancement are discussed. Our work enables automatic focus+context visualization, where the viewpoint information is taken into account. We therefore point out some previous focus+context approaches. Then, various rendering techniques for *levels of sparseness* specification are reviewed. Finally former work on incorporating cut-away concepts in visualization is mentioned.

Feature Classification: A typical feature classification in volume visualization is done through *transfer function* specification [18]. The transfer function with density as single input parameter is also denoted as one-dimensional transfer function. In recent years the idea of multi-dimensional transfer functions has been introduced. The multi-dimensional concept incorporates first and second derivatives of the density into the transfer function design [11], [17]. It is possible to assign optical

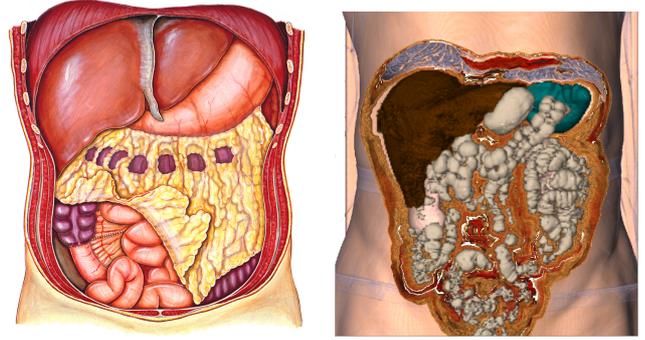


Fig. 1. Comparison between an artistic medical illustration of the abdomen (left) and our method (right).

properties based on gradient and curvature values, so for example object boundaries are classified differently than homogeneous regions. Taking into consideration first and second derivatives enables the semi-automatic generation of transfer functions [16]. An interesting approach was presented by Hauser and Mlejnek [9] for multi-dimensional 3D flow data. They use the *Degree-of-Interest* function that maps the user interest to optical properties. These concepts, however, define the representation globally, i.e., the visibility of important features is not guaranteed. The lack of adaptation of optical properties to the viewpoint settings is the main drawback of view-independent classifications.

Focus+Context Rendering: Visualization tasks frequently use the focus+context metaphor to clearly differentiate very relevant information from the context. Viewpoint-dependent distortion of three-dimensional data [3], for example, highlights data by dedicating more display 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* into the volume rendering pipeline [29]. The optical properties are changing according to the distance to the focal point. Using this technique several expressive focus+context effects can be achieved. Focus+context approaches use viewpoint-independent classification and therefore they have the same limitations as the feature classification methods discussed above.

Gaze-directed volume rendering [19] was an early approach in volume visualization, where the observer's viewing direction was taken into consideration. The motivation in this case was to increase the rendering performance instead of increasing the visual information. The volume dataset is rendered in different resolutions. According to the viewing direction only the focal region is represented in full resolution, and the other parts are rendered in lower resolution.

Sparse Representation: The graphics community has been inspired by artists to represent features sparsely in order to exploit the human imagination. The display of contours is a popular method to thinly represent context information in volume visualization [4], [24]. Outlines are often sufficient to roughly understand the shape and can be combined with other rendering techniques such as direct volume rendering or maximum intensity projection [10]. To make a contour representation more expressive *suggestive contours* can be used [5]. The suggestive contours technique combines contours rendered from a particular viewpoint with contours from other viewpoints close to the current view. 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 [26]. This is up to a certain degree similar to the combination of curvature-directed strokes with iso-surface rendering [14]. This approach was proposed for rendering structures that are nested within other objects. The interior structures are rendered fully opaque, while the enclosing objects are represented by a set of curvature-directed lines. Using stippling techniques for volume visualization is another example of inspiration from traditional illustration [21]. The visibility of interior structures can also be modified by dynamic changes in transparency of the outer shape. Recent work proposes to map transparency to the level of specular highlight [2]. This allows to see *inside* the volume in the areas of highlights. Dynamic transparency is also used in the 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 [13]. In volume visualization this technique is also known as *volume cutting* [23]. Automatic generation of cut-away images has been already researched in computer graphics [6], [7]. Straka et al. [25] are applying a cut-away technique for CT-Angiography. For visualizing complex dynamical systems, streamarrows were proposed by Löffelmann et al. [20]. They use arrows as a basic element for cutting away part of the stream surface. This allows to see through the surface and perceive other surfaces or structures behind.

Importance-driven volume rendering (IDVR) has been firstly introduced in our previous work [28]. We have presented a generalized model for view-dependent focus+context tasks in volume visualization. This paper extends the discussion on the key elements of the model: *object importance*, *levels of sparseness*, and *importance compositing*. Furthermore suggestions on feature defini-

tion are discussed to overcome the limited applicability to pre-segmented data. IDVR turns out to be helpful not only for changing viewpoints, but also in case of dynamic features.

III. 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 in the final image.

Interesting structures within the volumetric data are denoted as *features* respectively *objects*. The specification can be done in many different ways, also depending on the type of the data. In medical visualization features are often classified as particular organs. Such objects are defined by a segmentation process. Another way of feature classification can be through a spatial relationship within the volume or eventually through the location with respect to another *feature*. In such a way it is for example possible to classify the vortex core of a hurricane. In the case of multi-dimensional volumetric data, features can be defined by specifying interesting ranges of values for each data dimension. There are many other ways to determine features. A detailed treatment of feature definition is outside the scope of this paper.

Traditionally features within the volume dataset are classified by optical properties such as color and opacity. We additionally assign another dimension to features, which describes their *importance*. Importance encodes which features are most interesting and have the highest priority to be clearly visible. Each feature is therefore weighted by a positive scalar value called *object importance*. During the rendering stage, the model evaluates the visibility of each feature according to its importance. If less important structures 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. All interesting structures are visible irrespective if they are covered 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

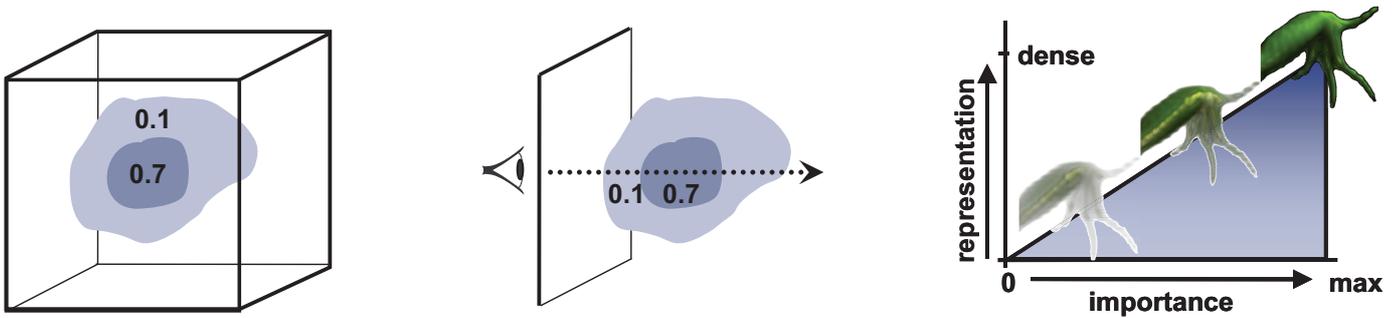


Fig. 2. Basic elements of importance-driven volume rendering: Volumetric features are classified by *importance* values (left). The volume is traversed in the *importance compositing* step (middle). Then *levels of sparseness* are chosen to enhance or suppress particular parts of the volume (right).

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 objects behind it. 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, denoted as *importance compositing*. This step evaluates the occlusion, takes the importance factor of each object into account and assigns to all objects particular *levels of sparseness*. The final synthesis results in images with maximal visual information with respect to the predefined object importance.

The relationship between the above mentioned components is depicted in Figure 2. *Importance compositing* is done similar to the direct volume rendering (DVR) approach. For each ray the compositing step evaluates object occlusions and assigns the corresponding *level of sparseness* to each object. Object importance is preserved in the sense that it is mapped to object visibility in the resulting image. This causes different rendering settings for the context object (importance value 0.1 in Figure 2) in the area of the image which is covered by the focus object (importance 0.7).

The difference between traditional volume rendering and importance-driven volume rendering is clearly visible in Figure 3. The goal is to emphasize the inner organ as focus object. 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.

IV. 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 three methods of importance compositing, which are inspired by compositing optical properties through ray casting of volume data.

A. Maximum Importance Projection

Maximum intensity projection (MIP) [22] is a simple and fast volume rendering approach. It is applicable for sparse data where important information has high intensity values such as 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 MIP 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,



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

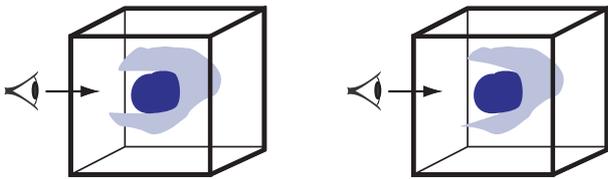


Fig. 4. Maximum Importance Projection. Illustration of cylindrical (left) and conical countersink (right).

i.e., fully transparent. With MImP structures are either rendered using the most dense representation or they are not rendered at all.

With MIP the spatial arrangement of structures is not readily apparent. MImP has a similar problem which we alleviate as follows: The image area, where the most important object is projected onto, is denoted as object *footprint*. With MImP the footprint is exactly the image region where only the focus 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 region 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 4 illustrates the difference between the cylindrical and conical MImP in 2D. The conical MImP is easily realized during ray traversal by changing the starting point for those rays that intersect the side faces of the clipping frustum. Figure 5 shows images to compare both approaches. The cylindrical MImP does not clearly show the spatial relationship between focus and context objects, i.e., the focus object appears in front of the context object. Conical MImP corrects this visual artifact.

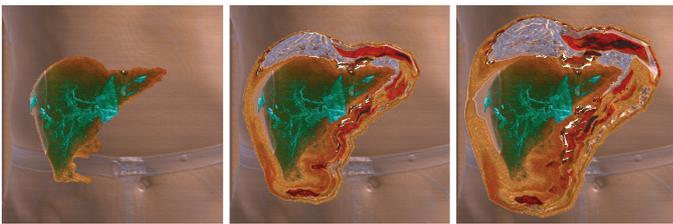


Fig. 5. Maximum importance projection (MImP). Cylindrical MImP (left) and conical MImP with different slope factors (middle and right).

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 object's *last hit* 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 object's last hit depth is reached. For the conical MImP we need to enlarge the footprint to build the conical shape. This is realized 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]. First e_{max} the highest depth value of the footprint is calculated. The starting points e_i of the rays that pierce the countersink are calculated from e_{max} , the slope s_c of the countersink, and distance d_i as shown in Equation 1:

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

d_i denotes the image space distance from pixel i to the footprint.

To correctly simulate the cut-out it is necessary to change the gradient vector of the ray entry points at the countersink geometry. 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* s_c of the countersink frustum.

B. 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. An object o has an importance value I_o . Ray r intersects 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 (AImC) does not completely remove the less important objects as with MImP. The sparseness factors are estimated according to the given importance. This allows a very sparse representation of the occluding object to see a rough structure of the shape and to clearly see the important object behind it. The importance compositing stage computes the levels of sparseness of all objects for every pixel in the final image. Levels of sparseness are computed using the object footprints. At each pixel position we perform a lookup to each object footprint. Object importance values of all objects that cover the current pixel position are summed up. The sparseness factor of each of these objects is estimated through division of their object importance by the evaluated sum (Equation 2).

The final image synthesis using AImC is an extension to traditional DVR. At each sample location during ray

traversal the level of sparseness additionally modulates the visibility of the sample. Similar to cylindrical MImP, AImC generates images where the spatial arrangement of structures is not readily apparent. In order to improve spatial perception we propose two methods to perform final importance-driven image synthesis using AImC, i.e., an image-space and an object-space approach.

Image-Space AImC: The object footprints introduce sharp transitions in levels of sparseness, which might reduce spatial perception. To improve the spatial perception we generate smooth transitions between different levels of sparseness. Before the levels of sparseness are computed for each object, we apply image processing operators to every footprint, i.e., dilation and averaging. As pixels in the footprint have a weight of one, the weights in the generated transition zone are smoothly decreasing to zero. The levels of sparseness estimation is analogous to Equation 2. For each pixel we compute the footprint-weighted importance sum of all contributing objects. The object importance is in this case always multiplied by the footprint value in the range of $[0, 1]$. Footprint values below one are part of the *transition* area between different levels of sparseness.

The image-space approach does not evaluate whether the sample of a suppressed context object is in front or behind the important object. The level of sparseness is constant for all samples of an object along a particular ray. This means that also part of the context object *behind* the focus object is suppressed.

Object-space AImC: To avoid suppression of context behind the focus object, we propose a more costly object-space approach. Using this approach only those samples of the context object are suppressed that are *in front* of the focus. In this case the level of sparseness is not constant for an object along a particular ray. The difference between the image-space and the object space approach is illustrated in Figure 6. The Figure shows that image-space AImC suppresses all context samples along a ray. The object space approach suppresses only the part of the context object that occludes the focus object.

The image synthesis of the object space approach is analogous to the conical MImP. In the case of the conical

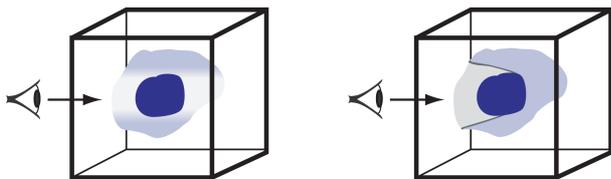


Fig. 6. Average importance compositing (AImC). The illustration depicts the difference between the image-space (left) and object space approach (right).

MImP the countersink geometry is used to estimate the starting position of the ray in order to perform the cut-out. In object-space AImC the countersink geometry defines the border between different levels of sparseness. The starting position of the ray is not changed. During the ray traversal in the final image synthesis, each sample location is evaluated whether it *belongs* to the countersink region or not. The context outside the countersink is depicted with a more dense representation and inside a more sparse form is chosen.

Results of each approach are shown in Figure 7. The images show the same dataset under different viewing angles. The top row illustrates the image-space approach and the bottom row the object-space approach.

The AImC approach preserves the *thickness* of the occluding part of the context object. This leads to different visibilities of the focus object under different viewing conditions. If the occluding context area is too thick, the focus object is not visible enough. In order to see the focus object properly the level of sparseness function has to be changed for the context object or the importance of the focus object has to be increased. The following Sub-section describes how to automatically overcome the problem a of varying thickness of the occluding object.

C. Visibility Preserving Importance Compositing

Visibility preserving importance compositing (VPImC) guarantees constant visibility of the focus object. Independent from the thickness of the context in front of the focus a constant fraction of visibility is reserved for the focus object. For example under some viewing angle the context object may be *thin* and samples can be represented more densely. Under a different viewing angle the context object may be *thicker* in front of the focus object. Therefore samples that belong to this area should be more transparent. This is illustrated in Figure 8. The suppression of the context varies according to the *thickness* of the context, therefore the visibility of the focus remains constant.

The values for levels of sparseness using VPImC are estimated in the same way as with AImC (Equation 2). In AImC the levels of sparseness are selected for each sample during ray-traversal. With VPImC we select an appropriate level of sparseness *after* the ray-traversal stage. The level of sparseness is determined as follows: The level of sparseness for the context object in front of the focus object shall be S_o . In VPImC the goal is to adjust the *average* accumulated opacity of the occluding region to be equal to the value S_o . Therefore the context part in front of the focus is rendered separately. All ray opacities of the occluding part are summed together to

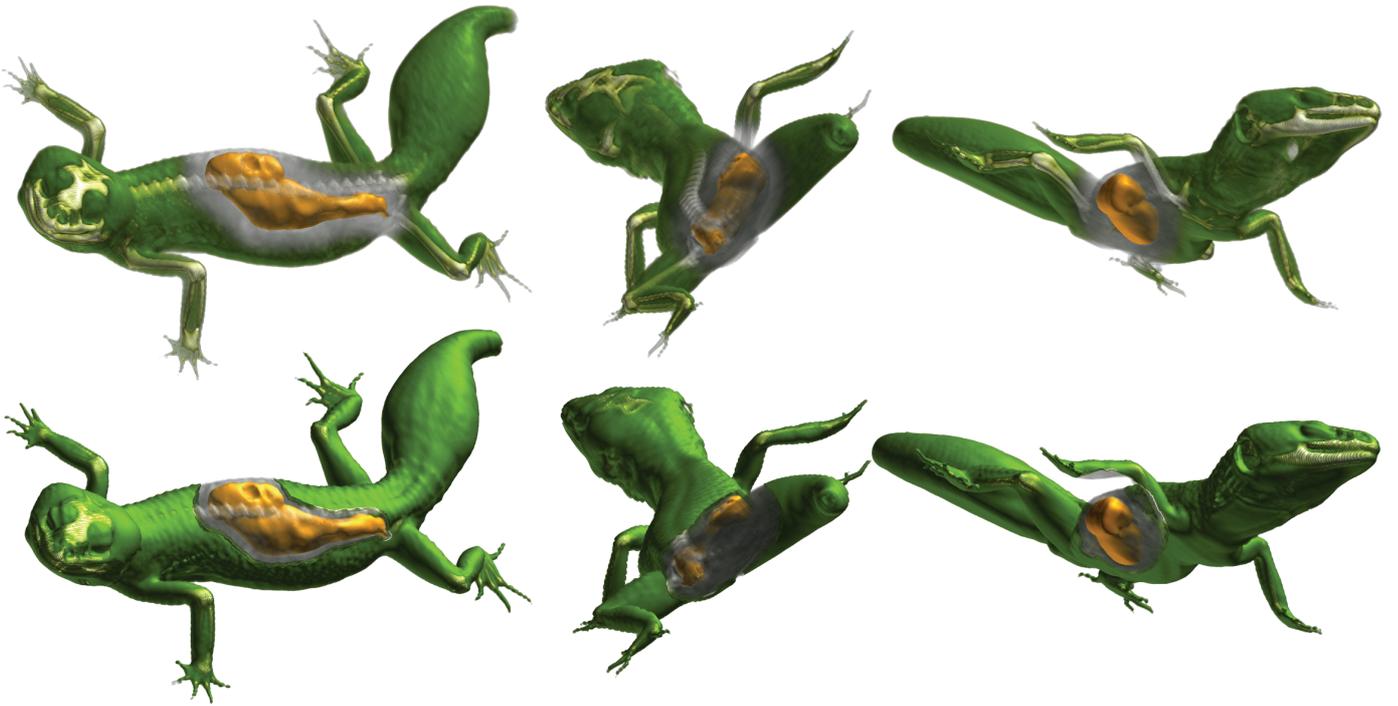


Fig. 7. Average importance compositing (AIcM) is shown under different viewpoint settings in combination with modulating optical properties. The upper row shows image based AIcM. The bottom row shows the object space approach with the same viewpoint settings.

compute the average per-ray opacity α_{avg} . To preserve the constant visibility of the focus object the average per-ray opacity has to be equal to S_o . Therefore for each ray, the accumulated opacity value is corrected. This is expressed by the Equation 3:

$$\alpha_{accum_new}(x) = \alpha_{accum}(x) \frac{S_o}{\alpha_{avg}} \quad (3)$$

where α_{accum_new} is the modified accumulated opacity of the suppressed context part, α_{accum} is the original accumulated opacity value, S_o is the level of sparseness value of the context object and α_{avg} is the average accumulated opacity value of the entire suppressed context part.

The separate rendering of the occluding part is done using two level volume rendering [10] (2IVR). Every

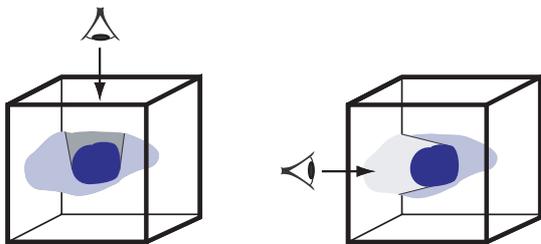


Fig. 8. Principle of visibility preserving importance compositing. Constant visibility allows a denser thin context (left) and requires a sparser thick context (right).

object or a part thereof (as in the case of the context object) is rendered separately in a *local compositing step*. Then the *visibility correction* is done for the occluding part. Finally a combination of both context parts takes place in the *global compositing step*. A more detailed discussion on 2IVR follows in Sub-section V-D.

Figure 9 shows the results of the compositing under different viewpoint settings. It offers a comparison to the importance compositing technique shown in Figure 7 (described in Section IV-B). Especially in the middle images of Figure 7 a large occluding context region considerably reduces the visibility of the focus object. This is not the case in Figure 9 (middle image.)

V. LEVELS OF SPARSENESS

Importance compositing determines for each part of an object its visibility in the rendered image. This is achieved by determining the appropriate level of sparseness. In the following four types of levels of sparseness are described. Three of those are depicted in Figure 10. The series of images illustrates how the context area in front of the focus object smoothly varies from a dense to a sparse representation.

A. Color and Opacity Modulation

A direct control of optical properties is the first approach to modify the visual prominence of a particular



Fig. 9. Visibility preserving importance compositing (VPImC). The dataset is shown under different viewpoint settings with constant visibility of the focus object. Focus visibility is independent from the thickness of the occluding context part.

feature. 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 through the level of sparseness. In general color is a very important visual cue in visualization. Highly saturated colors attract the observer's attention 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. 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 10 (top).

B. Screen-Door Transparency

Screen-door transparency is a well-known strategy to simulate transparency. The appearance of an object behind another semi-transparent object is simulated with a screen-door as follows: A screen-door consists of a wire-mesh and holes inbetween. The wires of the mesh depict the first object, while the second object is visible through the holes. From a certain distance holes and wires blend together to produce a semi-transparent impression. We use an analogous 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 arrangement of visible voxels is uniform and is forming a 3D wireframe structure. The impact of increasing sparseness is shown in Figure 10 (middle).

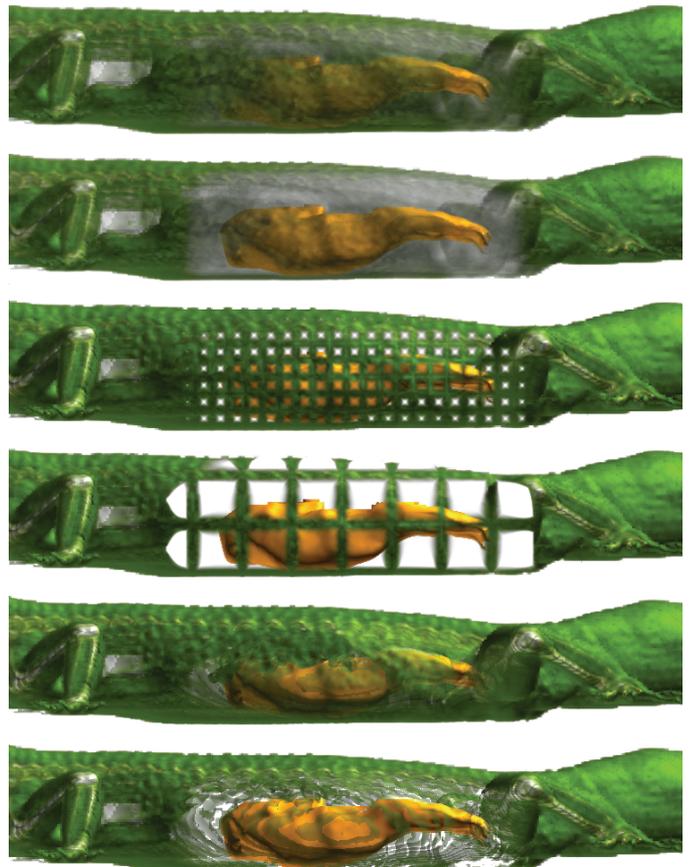


Fig. 10. Changing levels of sparseness. Top two rows: opacity modulation and color saturation modulation. Middle two rows: screen-door transparency. Bottom two rows: volume thinning. Images display levels of sparseness for the body with factors 0.75 and 0.25.

C. 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 con-



Fig. 11. Sparse and dense rendering styles: The occluding context object is rendered using summation (left), illustrative contour enhancement (middle), and maximum intensity projection (right).

tinuously 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 on strong iso-surfaces remain. Figure 10 (bottom) illustrates visibility reduction through volume thinning.

D. Sparse and Dense Rendering Styles

The previous levels of sparseness techniques describe how to enhance/suppress the visual representation of a particular object. The sparseness function smoothly varies from the most dense to the most sparse representation.

Another approach is to assign different *rendering techniques* as different levels of sparseness. For example the dense representation is achieved by direct volume rendering and the sparse one by illustrative contour rendering. In the case of object representations, combination is done via compositing. The combination of

different rendering techniques is achieved through two level volume rendering [10]. The difference between levels of sparseness based on object representations and rendering styles is shown in Figure 12.

Two level volume rendering (2IVR) is a technique to combine different volume rendering techniques. Well-known rendering techniques are direct volume rendering (DVR), MIP, summation (similar to X-Ray imaging), or illustrative rendering with contour enhancement [4], [24]. 2IVR renders each object within a volume with a different technique and composites the optical properties in a *local compositing step*. Each ray is partitioned by the intersecting objects into sub-rays. Local compositing is done for each sub-ray according to the rendering technique chosen for the respective object. The result of an entire ray is computed in a *global compositing step* which combines the results of the individual sub-rays.

Results of using different rendering techniques as levels of sparseness are shown in Figure 11. Where no occlusion occurs the context information is rendered using standard DVR. In the case of occlusion of the inner structure a different sparse rendering technique is applied. The images show the application of summation, contour enhancement, and maximum intensity projection for the local compositing. Global compositing is done by using again DVR.

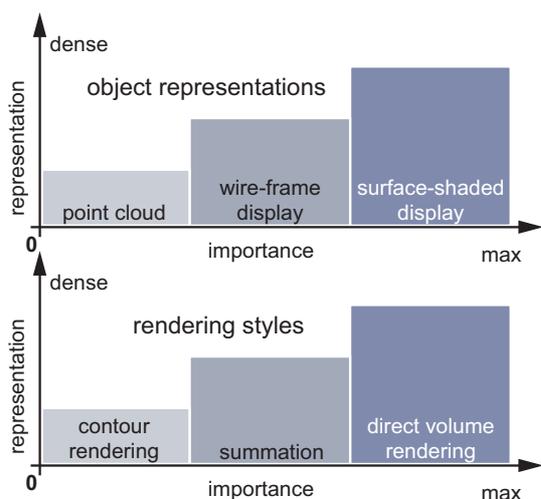


Fig. 12. Two types of *levels of sparseness*: based on object representations (top) and based on different rendering techniques (bottom).

VI. RESULTS

We show results of our method on three 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 third dataset is a time-varying simulation of the hurricane Isabel. It is a three-dimensional flow dataset with multiple simulated properties including cloud moisture, precipitation, pressure, and temperature. This dataset is not pre-segmented, only the position of the hurricane eye

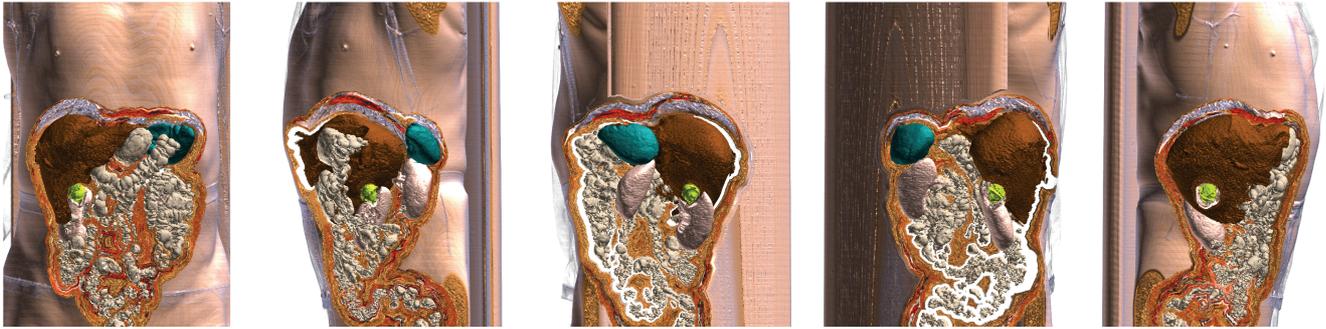


Fig. 13. The Monster Study dataset rendered using conical MImP. The highest importance is assigned to the tumor object (yellow). The organs of the abdomen are assigned a medium importance. The rest of the dataset has the lowest importance.

in object space is predefined. The simulation consists of 48 time steps.

Figure 13 shows the conical MImP of multiple abdomen organs from different viewpoints. The liver, spleen, kidneys, and intestine have the same importance value. The tumor (in yellow), located between kidney and liver, has the highest importance value. The rest of the body is of lower importance than any of the mentioned objects. The organs of the abdomen have the same importance and therefore they do not *cut away* each other. The highest importance value is assigned to the tumor and therefore everything in front of the tumor is cut-away. MImP allows to visualize the most important information, i.e., the tumor, its shape and spatial position in relationship to other objects. In contrast to traditional approaches the occlusion problem is solved automatically.

Another example of a conical MImP is shown in Figure 14. It shows 6 out of 48 time steps of the hurricane Isabel data. In this case a different way of feature classification was chosen. The important feature is the position of the eye of the hurricane. At the eye position a proxy cylinder is placed and everything inside the cylinder has higher importance than the rest of the data. The cylinder footprint is the basis for the countersink geometry.

This example also shows how to combine multiple scalar volumes using importance-driven volume rendering. The focus object is defined as the group of voxels inside the cylinder around the hurricane eye. Inside the cylinder the total precipitation mixing ratio is visualized. Thanks to the cut-away view it is possible to have a clear view at this property close to the eye of the hurricane. Outside the cylinder is the context area where the total cloud moisture is visualized. This time-dependent dataset also shows that the important feature can change its position and shape over time. Importance-driven volume

rendering guarantees to visualize the important feature irrespective of viewpoint and feature position and shape.

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

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

VII. SUMMARY AND CONCLUSIONS

In this paper we have investigated importance-driven volume rendering as a view and feature dependent approach 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.

Importance compositing defines how the occluding context information should be visualized. It can be simply cut away (MImP), or displayed using sparse representations. This approach can preserve either the thickness of the context object (AImC) or the visibility of the focus object (VPImC).

We have discussed four schemes for levels of sparseness. Levels of sparseness control 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

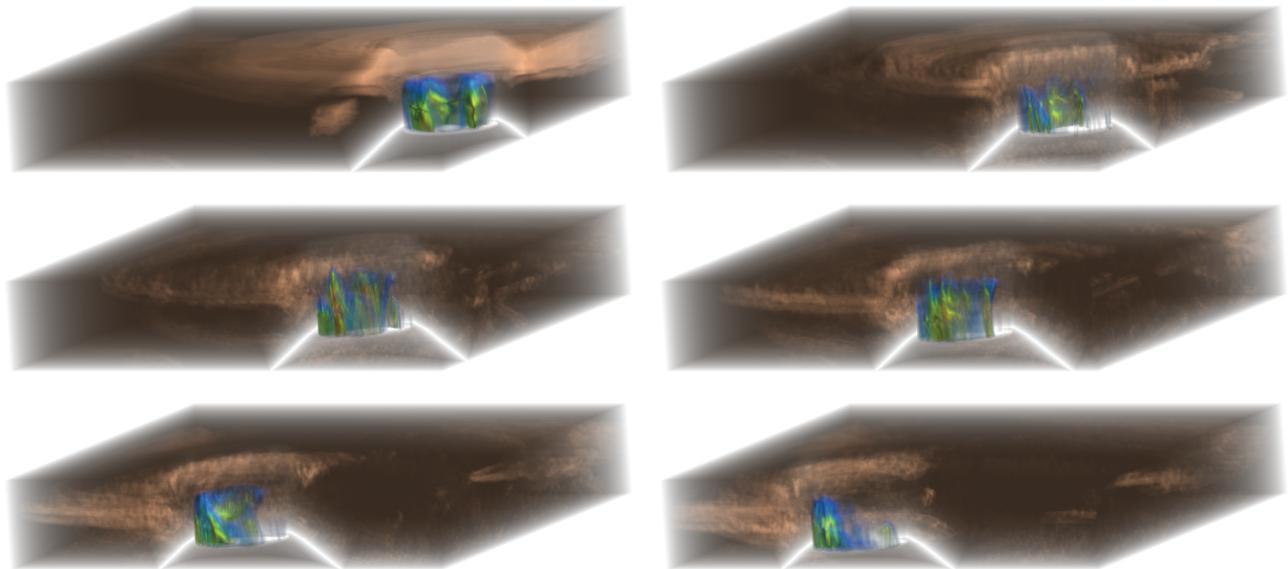


Fig. 14. Visualizing the simulation data of hurricane Isabel. Two different properties are visualized: total cloud and total precipitation mixing ratios. The interesting feature is the precipitation close to the eye of the hurricane. The context feature is the cloud mixing ratio. Images are showing 6 out of 48 time steps from left to right, top to bottom.

distributed uniformly over the whole volume, or the first- and second-order derivatives can be used for visibility distribution.

Levels of sparseness specify transitions of data representation from most dense to most sparse ones. Another approach defines levels of sparseness through different rendering techniques.

VIII. FUTURE WORK

The paper opens multiple opportunities for possible research areas. 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.

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 often 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 direction.

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. The parts with high first derivatives

values can then be considered as more important and a dense representation is chosen there.

The conical MImP and other object-space importance compositing approaches are implementing the cut-out illustration technique to improve perception of the spatial relationships. More elaborate approaches for intelligent automatic cut-out generations need to be researched. In cut-away views sometimes the borderline of the cut-out 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 contained in the data set. How to estimate viewpoint entropy and how to automatically determine optimal viewpoints [27] is another, not yet researched, area in volume visualization.

IX. 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 *Monster Study* dataset is courtesy of Tiani Medgraph and the *Leopard Gecko* is courtesy of University of the Veterinary Medicine Vienna. The hurricane Isabel data is courtesy of the National Center for Atmospheric Research in the United States. The medical illustration of the human abdomen is courtesy of Howell MediGraphics [12].

REFERENCES

- [1] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344–371, 1986.
- [2] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. Illustrative context-preserving volume rendering. Technical report, Institute of Computer Graphics and Algorithms, Vienna University of Technology, 2005.
- [3] M. S. T. Cappendale, 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.
- [4] 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.
- [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. Feiner and D. Seligmann. Cutaways and ghosting: Satisfying visibility constraints in dynamic 3d illustrations. *Visual Computer: International Journal of Computer Graphics*, 8(5-6):292–302, 1992.
- [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. I. Bischl, and M. E. Gröller. Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):242–252, 2001.
- [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, <http://www.medigraphics.com/>, 2004.
- [13] K. Hulsey technical illustration, <http://www.khulsey.com/>, 2004.
- [14] V. Interrante, H. Fuchs, and S. Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *Proceedings of IEEE Visualization '96*, pages 211–218, 1996.
- [15] J-Vision medical workstation, <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 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, 1988.
- [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] A. Lu, C. Morris, D. Ebert, P. Rheingans, and C. Hansen. Non-photorealistic volume rendering using stippling techniques. In *Proceedings of IEEE Visualization '02*, pages 211–218, 2002.
- [22] L. Mroz, H. Hauser, and M. E. Gröller. Interactive high-quality maximum intensity projection. In *Proceedings of EUROGRAPHICS '00*, pages 341–350, 2000.
- [23] 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.
- [24] P. Rheingans and D. Ebert. Volume illustration: Nonphotorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):253–264, 2001.
- [25] M. Straka, M. Červeňanský, A. La Cruz, A. Köchl, M. Šrámek, Eduard Gröller, and Dominik Fleischmann. The VesselGlyph: Focus & context visualization in CT-angiography. In *Proceedings of IEEE Visualization '04*, pages 385–392, 2004.
- [26] S. Treavett and M. Chen. Pen-and-ink rendering in volume visualisation. In *Proceedings of IEEE Visualization '00*, pages 203–210, 2000.
- [27] P. P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Viewpoint selection using viewpoint entropy. In *Proceedings of VMV '01*, pages 273–280, 2001.
- [28] I. Viola, A. Kanitsar, and M. E. Gröller. Importance-driven volume rendering. In *Proceedings of IEEE Visualization '04*, pages 139–145, 2004.
- [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.



Ivan Viola graduated in 2002 from the Vienna University of Technology, Austria, as a Dipl.-Ing. (MSc) in the field of computer graphics and visualization. Since then he is a PhD student and research associate in the ADAPT research project in the field of medical visualization and real-time volume visualization. His interests are efficient visualization in terms of quality, performance, and visual information.



Armin Kanitsar is key developer in the 3D department of Tiani Medgraph and research associate at the Institute of Computer Graphics and Algorithms (ICGA), Vienna University of Technology. He received his PhD in 2004 for his thesis 'Curved Planar Reformation for Vessel Visualization'. His main research interests include volume visualization, medical visualization, and image processing.



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