

Instant Volume Visualization using Maximum Intensity Difference Accumulation

Stefan Bruckner and M. Eduard Gröller

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

Abstract

It has long been recognized that transfer function setup for Direct Volume Rendering (DVR) is crucial to its usability. However, the task of finding an appropriate transfer function is complex and time-consuming even for experts. Thus, in many practical applications simpler techniques which do not rely on complex transfer functions are employed. One common example is Maximum Intensity Projection (MIP) which depicts the maximum value along each viewing ray. In this paper, we introduce Maximum Intensity Difference Accumulation (MIDA), a new approach which combines the advantages of DVR and MIP. Like MIP, MIDA exploits common data characteristics and hence does not require complex transfer functions to generate good visualization results. It does, however, feature occlusion and shape cues similar to DVR. Furthermore, we show that MIDA – in addition to being a useful technique in its own right – can be used to smoothly transition between DVR and MIP in an intuitive manner. MIDA can be easily implemented using volume raycasting and achieves real-time performance on current graphics hardware.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms

1. Introduction

Direct Volume Rendering (DVR) and Maximum Intensity Projection (MIP) are two of the most common methods for the visualization of volumetric data. DVR employs a physically-motivated absorption-plus-emission optical model and frequently utilizes gradient-based shading to emphasize surface structures. The basis of MIP, on the other hand, is the assumption that the most relevant structures for the investigation at hand have higher intensity values. In practice, this is achieved through special scanning protocols or the administration of contrast agents. Common examples include CT and MRI angiography as well as Positron Emission Tomography (PET), but also different microscopic imaging modalities where structures of interest are frequently highlighted using fluorescent marker proteins.

One of the biggest advantages of MIP over DVR is that it does not require the specification of complex transfer functions to generate good visualization results. A major disadvantage is, however, that due to the order-independency of the maximum operator, spatial context is lost. This paper in-

troduces a novel method which aims to combine the advantages of DVR and MIP. Our new approach is able to generate meaningful visualizations using a class of very simple linear transfer functions specified using standard window/level controls.

The remainder of this paper is structured as follows: In Section 2 we discuss the foundations of DVR and MIP and identify the drawbacks of these common methods. Section 3 reviews related work. Our novel rendering technique is detailed in Section 4. In Section 5 we present an approach to smoothly transition between DVR and MIP using our new method as an intermediate step. Results are presented and discussed in Section 6. The paper is concluded in Section 7.

2. Background

MIP works by traversing all viewing rays and finding the maximum data value along each of them. This maximum is then mapped to a color value and displayed to the user. In most cases this mapping process is simply a

linear transformation of data values to pixel intensity. As only a single value is displayed along each ray, MIP images lack depth information which can lead to visual ambiguities. Two approaches have been proposed to eliminate this drawback: Depth-Shaded Maximum Intensity Projection (DMIP) and Local Maximum Intensity Projection (LMIP). In DMIP [HMS95], the data value is additionally weighted by a depth-dependent term. This reduces the likelihood of high data values being projected if they are located far away from the image plane. While this approach can help to regain spatial context, it may also hide certain high-intensity regions. For LMIP [SSN*98], the first local maximum which is above a user-defined threshold is depicted. If no value above the threshold is found along a ray, the global maximum along the ray is used for the pixel. This approach also adds spatial information at the cost of introducing an additional parameter which can greatly affect the visualization.

DVR commonly employs a simplified model of light propagation in participating media [Max95]. It only accounts for emission and absorption of light but neglects scattering effects. Emission and absorption properties are specified using a transfer function which assigns color and opacity to each data value. The final color along each viewing ray is then determined through accumulation of colors and opacities at constant intervals – an approximative solution of the volume rendering integral. In order to enhance the appearance of surface structures in the volume, the vector of first partial derivatives along the three major axes – the gradient – can be used to evaluate a surface shading model. The color at each sample point along a ray is additionally modulated by the result of this evaluation. Surface-based shading generally helps to enhance fine details in areas with high contrast, but can lead to artifacts in nearly homogeneous areas. A major problem of DVR is the specification of an appropriate transfer function since assigning high opacity to a certain data range may occlude other structures of interest [PLB*01]. Thus, MIP is often the method of choice as it does not require additional parameter tuning even if DVR could lead to additional insight. In particular, shading information in DVR can help to interpret certain structures such as blood vessels.

In this paper, we propose a new rendering technique which aims to fuse the complementary advantages of DVR and MIP. Specifically, we want to preserve the practically parameterless nature of MIP and combine it with the added spatial context of DVR provided by accumulation and shading.

3. Related Work

DVR was introduced by Drebin et al. [DCH88] and Levoy [Lev88]. Since then, much work has focused on the task of improving the specification of transfer functions. Kniss et al. [KKH02] use a two-dimensional transfer function based on scalar value and gradient magnitude to effec-

tively extract specific material boundaries and convey subtle surface properties. Correa and Ma [CM08] propose the use of transfer functions based on the relative size of features to improve classification. In order to cope with the complexity of transfer function specification, several automatic and semi-automatic approaches for their generation have been proposed. Bergman et al. [BRT95] present an interactive approach for guiding the process of colormap selection. The contour spectrum, introduced by Bajaj et al. [BPS97], assists iso-value selection by presenting the user with 2D plots of several properties computed over the data range. He et al. [HHKP96] treat the search for a good transfer function as an optimization problem and employ stochastic techniques for this purpose. Marks et al. [MAB*97] sample the vast parameter space to find a set of input-parameter vectors that optimally disperses the output-value vectors, organizing the resulting graphics for easy and intuitive browsing by the user. Kindlmann and Durkin [KD98] employ histogram volumes of the data value and its first and second directional derivatives along the gradient direction to find a transfer function which makes boundaries more visible. Tzeng et al. [TLM05] propose the use of machine learning to find an optimal classification for user-specified regions of interest defined using a painting metaphor. The drawback of automatic approaches is that they require substantial pre-processing which limits their use in interactive settings where a quick exploration of the data is required.

Other approaches rely on properties of the data acquisition process to quickly generate meaningful visualizations. MIP, first introduced by Wallis et al. [WMAK89] as Maximum Activity Projection for Positron Emission Tomography (PET) data, emphasizes high-intensity data values. In addition to several optimization techniques [MKG00] it has been attempted to re-introduce missing spatial context information to MIP. Heidrich et al. [HMS95] proposed to use a depth-based weighting of the data values along a ray. Sato et al. [SSN*98] project the first local maximum instead of the global maximum. Hauser et al. [HMBG01] fuse DVR and MIP rendering for a single data set. Based on a pre-classification different compositing strategies are used for structures in a data set. Similarly, Straka et al. [SCC*04] employ a combination of rendering techniques for the visualization of vascular structures. Mora and Ebert [ME04] experiment with different order-independent operators such as maximum-projection and summation for the generation of overview images.

The area of illustrative visualization has investigated techniques to enhance visual comprehension based on concepts common in traditional illustration. Rheingans and Ebert [RE01] present several illustrative techniques which enhance features and add depth and orientation cues. They also propose to locally apply these methods for regional enhancement. Csébfalvi et al. [CMH*01] present a non-photorealistic technique to quickly generate contour-based overview images. Viola et al. [VKG05] introduce




Figure 1: Full body CT angiography rendered using (a) DVR, (b) MIDA, and (c) MIP. Data set courtesy of the OsiriX Foundation (<http://www.osirix-viewer.com>).

importance-driven volume rendering which generates cut-away views based on the importance of pre-classified objects. Bruckner et al. [BGKG06] present an illustrative volume rendering technique inspired by ghosted views. Rezk-Salama and Kolb [RSK06] introduce opacity peeling for the extraction of feature layers to enable the visualization of structures which are difficult to classify using transfer functions. Malik et al. [MMG07] extend this work with a more detailed analysis of ray profiles. Locally adaptive volume rendering, presented by Marchesin et al. [MDM07], attempts to reduce occlusion by dynamically adapting the opacity of sample contributions.

4. Maximum Intensity Difference Accumulation

The basic idea behind our approach is to alter the accumulation behavior of DVR to exhibit characteristics similar to MIP. Conventionally in DVR, when tracing viewing rays starting from the eye, the accumulated opacity is a monotonically growing function. This means that structures located behind thick non-transparent regions tend to have less influence on the final image. Because of this, a simple transfer function, such as a linear ramp, frequently causes structures of interest to be immersed in "fog", i.e., they are occluded by irrelevant material. While more complex transfer functions can be employed to remedy this problem, their spec-

ification is significantly more time-consuming. MIP, on the other hand, completely disregards such occlusion relationships. We want to adapt the behavior of DVR to prevent local maxima from becoming completely occluded while still preserving opacity-based accumulation.

We assume a continuous scalar-valued volumetric function $f(P)$ of normalized data values in the range $[0, 1]$. At the i -th sample location P_i along a viewing ray, f_{P_i} denotes the data value at location P_i and f_{max_i} is the current maximum value along the ray. Front-to-back traversal is used. We use $c(f_{P_i})$ and $\alpha(f_{P_i})$ to denote the color and opacity, respectively, of the sample value as classified by a transfer function. The accumulated color and opacity at the i -th sample position along the ray are denoted by c_i and α_i , c_0 and α_0 are initialized to zero.

We are interested in regions where the maximum along the ray changes. Specifically, when the maximum changes from a low to a high value, the corresponding sample should have more influence on the final image compared to the case where the difference is only small. We use δ_i to classify this change at every sample location:

$$\delta_i = \begin{cases} f_{P_i} - f_{max_i} & \text{if } f_{P_i} > f_{max_i} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$




Figure 2: Typical ray profiles for (a) DVR, (b) MIDA, and (c) MIP.

Whenever a new maximum is encountered while traversing the ray, δ_i is nonzero. These are the cases where we want to override occlusion relationships. For this purpose, the previously accumulated color c_{i-1} and opacity α_{i-1} are weighted by a factor $\beta_i = 1 - \delta_i$. For our new model, which we term Maximum Intensity Difference Accumulation (MIDA), the accumulated color c_i and opacity α_i for the i -th sample are then:

$$\begin{aligned} c_i &= \beta_i c_{i-1} + (1 - \beta_i \alpha_{i-1}) \alpha(f_{P_i}) c(f_{P_i}) \\ \alpha_i &= \beta_i \alpha_{i-1} + (1 - \beta_i \alpha_{i-1}) \alpha(f_{P_i}) \end{aligned} \quad (2)$$

Equation 2 only differs from standard DVR compositing by the additional weighting of c_{i-1} and α_{i-1} with β_i . One way to interpret the modulation of previously accumulated color and opacity performed in MIDA is a particular importance function which assigns highest prominence to local maxima. In particular, maxima which occur in a rather discontinuous manner cause more modulation than smoothly increasing ray profiles.

Figure 1 shows images generated using (a) DVR, (b) MIDA, and (c) MIP. Figure 2 depicts a typical ray profile for each of these techniques. A linear mapping of data values to grayscale intensities and (for DVR and MIDA) opacities is used. In the case of DVR, opacity is accumulated quickly resulting in a low overall intensity which manifests itself in the corresponding rendition as dark fog. For MIP, the intensity value along the ray is always equal to the current maximum. Using MIDA, due to the modulation of already accumulated intensities and opacities when a new maximum value is encountered, the intensity profile closely mimics the behavior of MIP. As visible by comparing Figure 1 (a) and (b), MIDA is able to immediately depict high-intensity features not visible in DVR without further transfer function modification. In contrast to the MIP result depicted in Figure 1 (c), however, MIDA features additional spatial cues due to accumulation which help to interpret feature locations.

4.1. Shading

In addition to accumulation, surface-based shading in DVR can provide important visual cues. It allows better judgement




Figure 3: Cranial MRI angiography rendered using (a) MIP without shading, (b) MIP with gradient-based shading, and (c) MIDA with gradient-based shading.

of shapes and fine details. MIP images, however, are usually devoid of this information. One reason for this is that in many cases the maximum along a ray may be located in a homogeneous area where surface shading produces bad results. Moreover, as the spatial location of the maximum between neighboring rays can vary considerably, lacking coherence between pixels can lead to disturbing artifacts. One major advantage of MIDA compared to common MIP is that the binary decision of which value to depict for each ray is replaced by accumulation. As the maximum normally does not change abruptly due to partial volume effects, samples in the boundary regions of high-intensity structures contribute to the ray color. This leads to improved coherency and allows us to perform surface-based shading as it is common in DVR without introducing the artifacts that would occur if it was applied to MIP. Figure 3 shows a comparison between (a) MIP, (b) MIP with shading applied, and (c) shaded MIDA. While shaded MIP fails to capture surface characteristics and introduces distracting artifacts, MIDA gives a good indication of the vascular shape. The magnified area shows how shading information can be helpful – a small aneurysm which is almost invisible in the MIP image is clearly recognizable using MIDA.

The gradient is a bad predictor for the surface normal orientation in nearly homogeneous regions due to the increased influence of noise. One approach to remedy this problem is to use the magnitude of the gradient vector to determine the degree of shading that is applied. In cases of low gradient magnitude the unshaded color is used, while higher gradient magnitudes result in an increased degree of shading. To achieve this, we linearly interpolate between the unshaded and the shaded color with an interpolation weight of $smoothstep(|\nabla f_{P_i}|, g_l, g_h)$ where $|\nabla f_{P_i}|$ is the gradient magnitude and g_l, g_h are lower and upper thresholds. The smoothstep function smoothly transitions from zero to one as the first argument varies between g_l and g_h and is commonly implemented as a cubic polynomial. For the thresholds we use the empirically determined values $g_l = 0.125$ and $g_h = 0.25$ which have delivered universally good results in all our ex-

periments and do not require user adjustment. This modulation is used in all images which feature shading throughout this paper.

4.2. Classification

A major advantage of MIDA, in contrast to DVR, is that it does not require complex transfer functions as the opacity profile is modified based on the difference between sample value and the current maximum. The result is an image which depicts the same essential features as a MIP rendering. We can therefore limit transfer function modification to brightness and contrast adjustment using common window/level controls which are almost universally incorporated in applications for the visualization of volume data. The user alters two parameters: the window w and the level l . The interval $[l - 0.5w, l + 0.5w]$ is then linearly mapped to the full range of grayscale intensities from black to white and opacities from zero to one. Alternatively, the user can choose other color maps based on domain conventions, but the opacity function does not require modification. Window/level adjustment is a routine task for domain experts. Additionally, methods for automatically finding good window/level settings are frequently employed at the time of acquisition and parameter values are commonly stored together with the data (e.g., in the DICOM format). However, MIDA already provides acceptable results for an even more restricted case: if the opacity is one and the color is white for all data values, MIDA – due to shading – depicts the essential features of the data set. Both, DVR and MIP would simply show a white image.

5. Combining DVR, MIP, and MIDA

Instead of advocating the complete replacement of DVR and MIP by our new method, we recognize that MIDA represents a good middle ground between these standard techniques. Some data sets can be visualized better with DVR, for others MIP is more suitable. In practice, some characteristics of both are desirable in many cases [FNH*06]. There-

fore one valuable approach is to use MIDA as a basis for a smooth transition between these different algorithms. We present a method where the user can smoothly make a visualization more DVR-like or MIP-like. As demonstrated, MIDA represents a hybrid between both methods and is therefore an ideal starting point. We introduce a new parameter $\gamma \in [-1, 1]$ defined as follows: For $\gamma = -1$, the rendering result is unmodified DVR, if $\gamma = 0$ only MIDA is used, and for $\gamma = 1$ the resulting image will be a MIP rendering. Values of $\gamma \in (-1, 0)$ will result in a smooth transition between DVR and MIDA, while values of $\gamma \in (0, 1)$ cause a transition from MIDA to MIP.

MIDA to DVR. As discussed in Section 4, MIDA introduces an additional modulation of the previously accumulated color and opacity along a viewing ray. If the modulation factor $\beta_i = 1$ for all samples along the ray, the result is DVR. For a smooth transition between MIDA and DVR, we simply modify β_i so that it approaches one for all samples when γ changes from MIDA to DVR in the following manner:

$$\beta_i = \begin{cases} 1 - \delta_i(1 + \gamma) & \text{if } \gamma < 0 \\ 1 - \delta_i & \text{otherwise} \end{cases} \quad (3)$$

Visually, this results in high intensity values smoothly fading out as γ is reduced. As the user moves γ towards DVR, more and more occlusion occurs.

MIDA to MIP. We could use a similar approach for the transition from MIDA to MIP, having β_i approach one only for samples where the maximum changes and letting it approach zero for all other cases. Essentially, regions where accumulation is performed would be "thinned" to the point where only one value along the ray is accumulated. However, shading would lead to problems in this case – the same artifacts as in Figure 3 (b) would arise as γ gets closer to one. Thus, we choose a different approach where the transition between MIDA and MIP is performed in image space. If $\gamma > 0$, we linearly interpolate between the accumulated MIDA color and opacity and the color and opacity of the maximum value after the ray has been traversed using γ as the interpolation weight. Interpolation is performed using opacity-weighted colors. Since MIDA and MIP images share the same basic characteristics, the major visual impacts of this transition are the gradual reduction of shading and the darkening of areas where much accumulation is performed.

6. Results and Discussion

The user can employ γ as a simple means for making a visualization result more DVR-like or more MIP-like. Instead of discretely switching between rendering methods, this allows a continuous manipulation of image characteristics on a clearly defined scale. The effects of modifying γ are easy to interpret: moving from MIDA ($\gamma = 0$) towards

DVR ($\gamma = -1$) increases occlusion – higher data values shine through less and less. Going from MIDA ($\gamma = 0$) to MIP ($\gamma = 1$), accumulation and shading are reduced.

In our experiments, we compared DVR, MIDA, and MIP applied to a number of different data sets. Figure 4 depicts an MRI angiography data set rendered using DVR, MIDA, and MIP. MIDA, like MIP, is able to depict the high-intensity vascular structures without further transfer function adjustment. However, MIDA provides additional spatial context as well as shape cues due to shading. In biomedical research, structures under investigation are commonly highlighted using fluorescent markers. MIP is therefore frequently employed for 3D visualization. DVR transfer functions are particularly difficult to find for these data sets, as there is no well-defined scale for the measured quantity. We experimented with several data sets from this field and MIDA helps to improve spatial comprehension compared to MIP. An example is shown in Figure 5, where a mouse embryo imaged using ultramicroscopy [DLS*07] is depicted. Due to shading, MIDA depicts vascular structures more clearly. Furthermore, as MIDA can generate meaningful images without requiring much transfer function tuning, it is well-suited for visualization tasks with strict time constraints. One example is the screening of luggage based on tomographic modalities such as CT. While conventional X-Ray is still the default modality for examining items in security-critical environments, it is often difficult to recognize potential threats due to its two-dimensional nature. Thus, in recent years CT-based screening has found increasing adoption in this area and is currently employed at many institutions such as airports and government buildings [YNG*07]. Figure 6 shows a CT scan of a backpack filled with various items – MIDA instantly allows the operator to inspect the contents and provides more information about the 3D structure of the individual items than MIP.

An existing DVR implementation can be extended to MIDA in a straight-forward manner by modifying compositing as described in Section 4. In terms of performance, the additional instructions required by MIDA typically can be neglected. However, as the opacity along a ray is no longer monotonically increasing, early ray termination is not possible. Also, in contrast to MIP, ray traversal can not be terminated when the current maximum is the highest intensity value in the data set. Despite the lack of these optimizations, current graphics hardware is still easily capable of rendering typical data sets at interactive frame rates. The average frame rates of our implementation measured on an NVidia GeForce 8800 GTX GPU for the standard UNC head data set ($256 \times 256 \times 224$) with a viewport size of 512×512 and an object sample distance of 1 were 9.9 (DVR), 9.1 (MIDA), and 18.2 (MIP). As MIDA performs only slightly worse than DVR, we have not further investigated potential acceleration techniques.




Figure 4: MRI scan rendered using (a) DVR, (b) MIDA, and (c) MIP. Data set courtesy of the OsiriX Foundation (<http://www.osirix-viewer.com>).




Figure 5: Ultramicroscopy of a mouse embryo rendered using (a) DVR, (b) MIDA, and (c) MIP. Data set courtesy of Dodt et al. [DLS*07].

7. Conclusions

In this paper, we introduced MIDA, a simple unifying extension of DVR and MIP which incorporates characteristics of both techniques. In contrast to DVR, MIDA does not rely on complex transfer functions but still features important spatial cues due to accumulation and shading. Furthermore, we presented an approach for smoothly interpolating between DVR, MIDA, and MIP. This allows users to enhance visualizations generated using DVR to include characteristics of MIP and vice versa. MIDA can be incorporated into existing volume rendering systems in a straight-forward manner. In experiments, our new technique has shown to achieve promising results for a wide range of different types of volume data.

Acknowledgements

We thank the anonymous reviewers for their valuable comments. The work presented in this publication was carried out as part of the **exvisation** project supported by the Austrian Science Fund (FWF) grant no. P18322.

References

- [BGKG06] BRUCKNER S., GRIMM S., KANITSAR A., GRÖLLER M. E.: Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization and Computer Graphics* 12, 6 (2006), 1559–1569.
- [BPS97] BAJAJ C. L., PASCUCCI V., SCHIKORE D. R.: The contour spectrum. In *Proceedings of IEEE Visualization 1997* (1997), pp. 167–173.
- [BRT95] BERGMAN L. D., ROGOWITZ B. E., TREINISH L. A.: A rule-based tool for assisting colormap selection. In *Proceedings of IEEE Visualization 1995* (1995), pp. 118–125.
- [CM08] CORREA C., MA K.-L.: Size-based transfer functions: A new volume exploration technique. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1380–1387.
- [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.
- [DCH88] DREBIN R. A., CARPENTER L., HANRAHAN P.: Volume rendering. In *Proceedings of ACM SIGGRAPH 1988* (1988), pp. 65–74.
- [DLS*07] DODT H.-U., LEISCHNER U., SCHIERLOH A., JÄHRLING N., MAUCH C. P., DEININGER K., DEUSSING




Figure 6: CT scan of a backpack filled with various items rendered using (a) DVR, (b) MIDA, and (c) MIP. Data set courtesy of Kevin Kreeger, Viatronix Inc. (<http://www.volvis.org>).

- J. M., EDER M., ZIEGLÄNSBERGER W., BECKER K.: Ultramicroscopy: Three-dimensional visualization of neuronal networks in the whole mouse brain. *Nature Methods* 4 (2007), 331–336.
- [FNH*06] FISHMAN E. K., NEY D. R., HEATH D. G., CORL F. M., HORTON K. M., JOHNSON P. T.: Volume rendering versus maximum intensity projection in CT angiography: What works best, when, and why. *Radiographics* 26, 3 (2006), 905–922.
- [HHKP96] HE T., HONG L., KAUFMAN A., PFISTER H.: Generation of transfer functions with stochastic search techniques. In *Proceedings of IEEE Visualization 1996* (1996), pp. 227–234.
- [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.
- [HMS95] HEIDRICH W., MCCOOL M., STEVENS J.: Interactive maximum projection volume rendering. In *Proceedings of IEEE Visualization 1995* (1995), pp. 11–18.
- [KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the IEEE Symposium on Volume Visualization 1998* (1998), pp. 79–86.
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 8, 3 (2002), 270–285.
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37.
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUML W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: a general approach to setting parameters for computer graphics and animation. In *Proceedings of ACM SIGGRAPH 1997* (1997), pp. 389–400.
- [Max95] MAX N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 1, 2 (1995), 99–108.
- [MDM07] MARCHESIN S., DISCHLER J.-M., MONGENET C.: Feature enhancement using locally adaptive volume rendering. In *Proceedings of the International Symposium on Volume Graphics 2007* (2007), pp. 41–48.
- [ME04] MORA B., EBERT D. S.: Instant volumetric understanding with order-independent volume rendering. *Computer Graphics Forum* 23, 9 (2004), 489–497.
- [MKG00] MROZ L., KÖNIG A., GRÖLLER M. E.: Maximum intensity projection at warp speed. *Computers & Graphics* 24, 3 (2000), 343–352.
- [MMG07] MALIK M. M., MÖLLER T., GRÖLLER M. E.: Feature peeling. In *Proceedings of Graphics Interface 2007* (2007), pp. 273–280.
- [PLB*01] PFISTER H., LORENSEN B., BAJAJ C., KINDLMANN G., SCHROEDER W., AVILA L. S., RAGHU K. M., MACHIRAJU R., LEE J.: The transfer function bake-off. *IEEE Computer Graphics and Applications* 21, 3 (2001), 16–22.
- [RE01] RHEINGANS P., EBERT D. S.: Volume illustration: Non-photorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (2001), 253–264.
- [RSK06] REZK-SALAMA C., KOLB A.: Opacity peeling for direct volume rendering. *Computer Graphics Forum* 25, 3 (2006), 597–606.
- [SCC*04] STRAKA M., Cervenansky M., CRUZ A. L., KÖCHL A., ŠRÁMEK M., GRÖLLER M. E., FLEISCHMANN D.: The VesselGlyph: Focus & context visualization in CT-angiography. In *Proceedings of IEEE Visualization 2004* (2004), pp. 385–392.
- [SSN*98] SATO Y., SHIRAGA N., NAKAJIMA S., TAMURA S., KIKINIS R.: Local maximum intensity projection (LMIP): A new rendering method for vascular visualization. *Journal of Computer Assisted Tomography* 22, 6 (1998), 912–917.
- [TLM05] TZENG F.-Y., LUM E. B., MA K.-L.: An intelligent system approach to higher-dimensional classification of volume data. *IEEE Transactions on Visualization and Computer Graphics* 11, 3 (2005), 273–284.
- [VKG05] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven feature enhancement in volume visualization. *IEEE Transactions on Visualization and Computer Graphics* 11, 4 (2005), 408–418.
- [WMAK89] WALLIS J. W., MILLER T. R., A. LERNER C., KLEERUP E. C.: Three-dimensional display in nuclear medicine. *IEEE Transactions on Medical Imaging* 8, 4 (1989), 297–303.
- [YNG*07] YING Z., NAIDU R., GUILBERT K., SCHAFFER D., CRAWFORD C. R.: Dual energy volumetric x-ray tomographic sensor for luggage screening. In *Proceedings of the IEEE Sensors Applications Symposium 2007* (2007), pp. 1–6.