

# Instant Volumetric Understanding with Order-Independent Volume Rendering

Benjamin Mora<sup>†\*</sup> and David S. Ebert<sup>‡</sup>  
<sup>†</sup>School of ECE    <sup>\*</sup>Envision Center  
Purdue University

## Abstract

*Rapid, visual understanding of volumetric datasets is a crucial outcome of a good volume rendering application, but few current volume rendering systems deliver this result. Our goal is to reduce the volumetric surfing that is required to understand volumetric features by conveying more information in fewer images. In order to achieve this goal, and in contrast with most current methods which still use optical models and alpha blending, our approach reintroduces the order-independent contribution of every sample along the ray in order to have an equiprobable visualization of all the volume samples. Therefore, we demonstrate how order-independent sampling can be suitable for fast volume understanding, show useful extensions to MIP and X-ray like renderings, and, finally, point out the special advantage of using stereo visualization in these models to circumvent the lack of depth cues.*

*Categories and Subject Descriptors:* 1.3.3 [Computer Graphics]: Picture/Image, Generation, 1.3.7 [Computer Graphics]: Three-Dimensional graphics and realism.

## 1. Introduction

An increased understanding of volumetric datasets is probably the greatest contribution of direct volume rendering (DVR), and the last two decades have seen a profusion of papers in this area. In early scientific visualization, maximum intensity projection (MIP) and X-ray projection were the most commonly used techniques. Today, shaded optically realistic renderings are more widely used and have many advantages, especially for visualization of ultrasounds datasets [HJC03] and virtual endoscopy [HMKB97]. However, DVR techniques are still not commonly used in the medical sector for several reasons, including usability, familiarity, demonstrated accuracy, reproducibility, and a lack of clinical proof of their value in conveying more information than traditional techniques. MIP/X-ray techniques have advantages over optically realistic DVR techniques in terms of their familiarity to the medical community and their ease of use. Therefore, we have developed improvements to MIP/X-ray renderings that harness these advantages and drastically reduce the spatial depth perception problem of these traditional DVR techniques. Besides introducing improvements to MIP/X-ray DVR renderings, we have studied the effectiveness of using stereo. Given the depth discrepancy of maximum values rendered with MIP, it is not clear whether or not stereo will be effective. However, our experience has shown that stereo drastically improves the understanding of spatial relationships within MIP and X-ray renderings.

## 2. Background

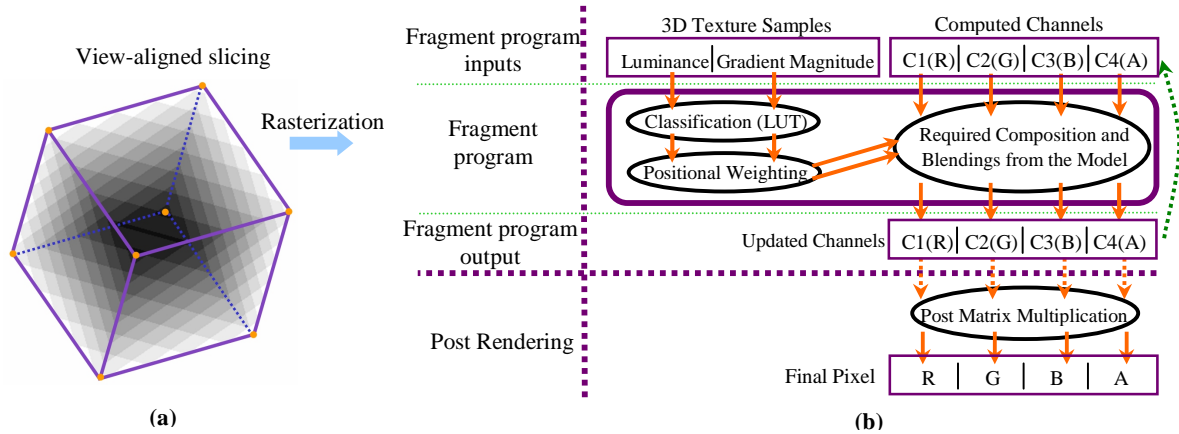
While the issue of realistic DVR techniques increasing perception of information over traditional techniques can be debated, it does highlight the questions of what makes an image informative and how can we improve the perception of our generated images.

Although there is probably no short answer, one can try to analyze the main differences between MIP/X-ray images and shaded DVR algorithms. Most often, these latter techniques [Lev88] use a rendering equation [KvH84] [Max95] that can be written as

$$I = \int_0^l C(s) \times \tau(s) \times \exp\left(-\int_0^s \tau(t) dt\right) ds \quad (1)$$

where  $I$  is the amount of light along the ray reaching the viewpoint,  $C(s)$  the amount of light emitted at the location  $s$ ,  $\tau(s)$  the extinction coefficient (i.e., the opacity at the sampling location), while the exponential represents the transparency (visibility) coefficient between the eye ( $s=0$ ) and  $s$ . Therefore, the accumulated colors and transparencies along the rays depend on two fundamental processes that are called classification and shading ( $\tau(s)$  and  $C(s)$ ) at the location  $s$ . Let us assume now that there are two average values (depending on the transfer function settings) for  $\tau(s)$  and  $C(s)$ , noted  $\tau_{mean}$  and  $C_{mean}$ , with  $0 < \tau_{mean}$ . Thus, the mean contribution,  $P(s)$ , for the  $s$  sampling location to the final pixel value is equal to:

$$\begin{aligned} P(s) &= C_{mean} \cdot \tau_{mean} \cdot \exp\left(-\int_0^s \tau_{mean} dt\right) \\ &= C_{mean} \cdot \tau_{mean} \cdot \exp(-\tau_{mean} \cdot s) \end{aligned} \quad (2)$$



**Figure 1:** Pipeline of our rendering system. View-aligned volume slicing at the vertex level (a) and fragment program (b)

The exponential part of (2) is important, and tells us that the probability of seeing a sample decreases exponentially with its distance to the eye, weighted by a constant depending on the transfer function (TF). Therefore, the closest samples are more prone to be seen, and visualization of further samples will be less probable, creating a need for navigating within the volume. Notice that a good transfer function for reducing navigation can be one that minimizes  $\tau_{mean}$ .

In reality, volumes are not so uniformly distributed, but in general, the occluded portions quickly increase according to the distance to the eye. A good example of this phenomenon is virtual endoscopy, where the viewpoint has to travel along the colon path to find possible polyps.

In contrast with traditional shaded methods, MIP/X-ray renderings can independently combine any sampled value in any order to obtain the final value. We will refer to these techniques as order-independent (volume) rendering (OIVR). Thus, without any *a priori* knowledge of the dataset, any volumetric feature has the same probability to be present in the final image, whether it is in front or in back. This property can theoretically reduce the navigation to only rotations. For instance, in virtual endoscopy, one might easily imagine the interest of such a method, since it would avoid a long navigation by instantly presenting all the volumetric features in the final image without occlusion.

However, it is still not clear how the human visual system perceives the spatial (volumetric) information of datasets rendered with OIVR techniques. While researchers have tried for many years to improve volumetric understanding of photorealistic renderings, it is only recently that research into improving non-photorealistic volume rendering has begun [CMH\*01, RE01, LM02, LME\*02, SE03]. However, both of these approaches use an alpha-blending model, which has been studied extensively in its effects on improving volume rendering. In contrast, the value of OIVR methods has not been widely studied. Most of the previous work has either mainly focused on the acceleration of OIVR [TL93, Ma193, MHG00, CSK03] or used OIVR methods when sample ordering was difficult to achieve [WKE02], instead of either improving or proposing

relevant models [BMDF02]. One of the best advances in this direction is the idea of fusing several rendering types, including both traditional DVR and OIVR techniques [HMBG01, HBH03].

Therefore, an appropriate use of OIVR models has strong potential, and improvements to OIVR techniques will be beneficial. In the next section, we describe several simple, new ways to improve OIVR algorithms. Our work has focused on how to convey more information from these well-known models through simple rendering improvements and on harnessing the great advantages of using stereo rendering with OIVR techniques.

### 3. Models and Implementation

This paper mainly focuses on new rendering models and can, therefore, be implemented with any volume rendering technique (ray-casting, splatting, etc...). However, since interactivity is desirable, we chose to use 3D textures, which slightly restricts the flexibility of our approach. Thus, Section 3.1 presents our texture-based implementation before introducing new ways to exploit OIVR.

#### 3.1. 3D Texture-based Implementation

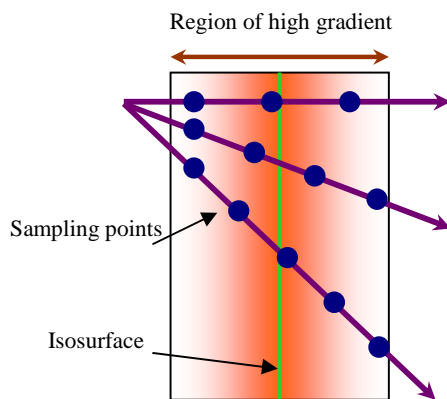
Our implementation (see Figure 1) uses a basic 3D texture-based volume rendering technique [WE98, EKE01], projecting view-aligned slices from a regular grid onto the screen (Figure 1a). The grid is represented as a 3D texture, where two 16-bit (12-bit significant) entities are stored per texture sample: the original dataset value and the gradient magnitude. This can be done by, for example, using NVidia HILO textures. The gradient magnitude is computed from a (smooth) 3D Sobel filter. Each slice-volume intersection is actually subdivided into a set of triangles where a 3D texture coordinate is assigned to every triangle vertex, corresponding to the vertex location inside the volume.

The OIVR model is then implemented at the fragment level by writing a (NV) fragment program (Figure 1b). A buffer was also used to render into an intermediate image. First, the fragment program accesses two tri-linearly

interpolated texture samples (original signal and its gradient) from the interpolated 3D texture coordinates passed to the fragment program. These samples are then mapped to values between 0 and 1 using a user-defined look-up table, resulting in two 4-component vectors that undergo an optional positional weighting (see Section 3.3) to emphasize local features. This step could occur before the look-up table step, but the post-look-up weighting seems to produce better results. These two vectors are then combined, according an OIVR model, with the four pixel's RGBA values to produce the new RGBA values. Notice that in order to get a read/write pixel access in the (NV) fragment program, the program binds the pbuffer in its own context. This method may generate undefined behaviors, but works fine with volume rendering (two pbuffers could be used instead). Also notice that a channel here is not a color, but a measure of an OIVR model (for example the *R* channel can be used to compute MIP, and the *G* channel can be used to compute the signal integral). However, for most of our examples, we use the same OIVR model, but with different transfer functions.

Finally after slicing, the four components undergo a post-matrix multiplication to obtain a final color image. This final multiplication enhances perception of the rendering by mapping the four results to different (user-defined) colors, utilizing the human color perception to increase understanding. This multiplication also handles both a possible brightness adjustment and *g*dividing channels by the number of slices whenever the *sum* operator is used.

Some implementation details of our interactive approach must also be mentioned. First, we use 32-bits floating-point buffers to render images without precision artifacts. For sum-based models, the final value needs to be divided by the total number of slices and, therefore, 8-bit precision is not sufficient to store the intermediate value. Second, since there is one fragment program per 4-component model, the number of possible fragment programs is equal to  $n^4$ , where  $n$  is the number of different models. Therefore, we only implement the most interesting models. Most of the time, the same model will be used for each of the four channels, which also allows a better parallelism inside the fragment



**Figure 2:** Summation around an isosurface. The greater the angle of incidence, the greater the number of accumulated samples in the gradient area (orange), resulting in a greater summed contribution for that region.

© The Eurographics Association and Blackwell Publishing 2004.

program than having different models.

### 3.2. OIVR Models

The four RGBA channels compute four OIVR models that are chosen from a set of 6 models. Notice that unlike optical models [KvH84, Max95], there is not a real mathematical justification for these models. These models simply produce different types of images that are useful for different purposes.

In order to be independent, the compositing of samples along the ray within one channel should be both commutative and associative. Thus, we choose the following two compositing functions: the *maximum* and the *sum* operators. These two functions are applied on three different entities derived from the two interpolated texture fragments: the original signal value, the gradient magnitude, and the original signal value times the gradient magnitude.

#### Original Signal

Applying either the *sum* operator or the *maximum* operator on the original signal value produces the famous X-ray [Mal93, TL93] or MIP models, respectively.

Research on X-ray like images has most often focused on Fourier volume rendering (FVR), where samples are linearly mapped according to their value. Here, our model can use any non-linear 1D transfer function, allowing the emphasis of specific regions of the signal range to produce better contrast in the images. In fact, we have found that the use of a sharp triangular (or Gaussian) peak in the transfer function can produce shaded, contour-enhanced, isosurface visualizations. The only condition is that the volume histogram should not contain any samples at the peak location to prevent non-isosurface values from being summed. Therefore, the pixel contributions are only from interpolated values. Moreover, the more parallel the ray is to the isosurface, the greater the number of sampled values around the isovalue (cf. Figure 2). This provides a contour enhancement similar in many ways to previous results [CMH\*01], but combined with OIVR visualization, and without the need to store a gradient vector in the 3D texture, which saves up to four times the video memory. However, the sampling rate needs to be increased compared to traditional texture-based algorithms.

#### Gradient Magnitude

The gradient magnitude is often used for classification in traditional volume rendering, since object boundaries are usually high-gradient regions. Thus, computing the gradient magnitude sum along a ray allows a better viewing of transitions between objects than a traditional X-ray sum. An interesting property shown in [TL93] is that this sum typically produces shaded renderings because a high gradient region usually revolves around the isosurface. Once again, the more parallel the rays are to the isosurface, the more samples within the gradient area (cf. Figure 2), and the greater the contribution to the final sum. However, summing the samples after applying a transfer function

mapping is, once again, more flexible than only computing the gradient magnitude sum along the ray.

We have also implemented a new model using the maximum of the gradient magnitude along the ray (GMIP). GMIP rendering produces compelling results as shown in Section 4.

### Original Signal x Gradient Magnitude


Recent research has shown that 2D transfer functions are very useful for classifying volume data [KKH02] from the gradient magnitude and the original signal. However, since storing a 2D table for 16 bits datasets requires too much memory, we use a slightly different approach that multiplies the two weights resulting from the sampled signal classification and the sampled gradient magnitude classification. Thus, this distinguishes between two object borders that could have the same gradient magnitude, but not the same intensity, and vice-versa.

### 3.3 Local Lighting of materials

Visualizing the entire volume when the interesting

volumetric features are located in small regions (local features) can sometimes decrease the understanding of those features. Therefore, our system allows the samples to undergo a positional weighting after the transfer function indexing. We refer to this as *local lighting* since the enhanced regions seem to be illuminated by that weighting, even if there is obviously no lighting model applied. Notice that this local weighting breaks our main idea that the contribution probability of every sample should not depend of the sample position.

The lighting weight of every sample decreases according to a distance that can be either the distance to a plane or the distance to a center point. Using the distance to a plane allows fast radar scanning of the volume, while a spherical distance allows quick focusing on the relevant regions within the volume. The sample weight is computed directly from the distance with either a Gaussian function or a windowed linear function. In both cases, some computational overhead is added to the fragment program, but the renderings are still interactive. Distance is computed inside the fragment program from the texture coordinate and, either the plane equation or the sphere center, stored in a constant register.



**Figure 3:** The UNC brain and the engine datasets rendered from our six different models. Rows differ from the measured entity at the sampling location (original signal, gradient magnitude, and product). Columns indicate the compositing function.

### 3.4. Stereo Order-Independent Volume Rendering

Using stereo visualization for OIVR is a key element of our system, and a new approach to OIVR rendering. We have implemented the *parallel axis asymmetric frustum perspective projection* stereo method. Notice that perspective rendering also contradicts once again our idea of equiprobable samples, since closest voxels will have a bigger projection. While using stereo volume visualization is not new, publications have only focused on photo-realistic models. The only previous stereo OIVR research, [PSL\*95], uses two static stereo MIP images to sketch vessel like structures. However, moving the viewpoint interactively was not possible, and there was no analysis of the effectiveness of stereo in MIP in general. This basic approach was insightful, since most OIVR research [Mal93, TL93, SSNK98, MHG00, SC02, CSK03] discussed the lack of occlusion in OIVR, without exploring stereo rendering as a solution. Normally, volume rotation was presented as a possible solution to circumvent depth occlusion. Variations of the MIP algorithm were also proposed in [SSNK98, SC02] to circumvent the lack of depth cues.


Therefore, the use of stereo in OIVR is probably more useful than in traditional renderings, which already generate depth cues, while there are no depth cues in OIVR. In fact, it is not obvious how the brain will interpret a stereo OIVR stimulus, whether it is from a MIP or an X-ray like image. While two projections are not enough to mathematically reconstruct a volume [CCF94], we have explored whether the visual system can solve the occlusion ambiguities and our results are discussed in the next section.

## 4. Results

For testing we have used the following five common volume datasets: the UNC CT head (256x256x225) and MRI brain (256x256x167) datasets, the CT engine (256x256x110) dataset, a MRI blood vessel (256x256x100) dataset, and a virtual endoscopy CT dataset (496x496x183). All renderings have been made on an Nvidia Quadro FX 3000 graphics card. For stereo, we used a 20 visible inch monitor with a refresh rate of 120 Hertz (60 Hertz per eye). Our implementation processes approximately 600 slices per second depending on both the dataset and the chosen model for a 512x512 image. Using our lighting technique decreases these rates by approximately 30%. Finally, images have been generated using a small sampling distance of between 0.25 and 0.5 (1.0 being the size of an isotropic voxel).

### 4.1 Testing the OIVR Models Separately (with affine TFs)

Our first test (see Figure 3) was to render the UNC brain and the engine datasets using our 6 OIVR models separately. Only one channel was used, producing black and white images. Quasi-linear transfer functions have been used to remove dataset noise that is usually present in the samples of low values. A quasi-linear transfer function can



**Figure 5:** Perspective MIP rendering of a blood vessel. While the whitest and thinnest central vessel seems to be in front of the larger blood vessel, it is in fact behind. Stereo allows making the difference, and the larger vessel is perceived as transparent at the intersection.

be defined as a partially affine function where  $f(x)$  linearly increases inside the range  $[a,b]$ ,  $f(x)=0$  for  $x<a$ , and  $f(x)=1$  otherwise.

Our results show that each model has the ability to display unique features, justifying the combination of the models to increase information in the final image. The difference is particularly visible between images generated from the gradient and images generated from the original signal. The gradient sum drastically improves the contrast over the signal sum by removing the homogenous regions of the volume. Contours and surfaces are now much more visible and appear shaded. Gradient MIP (GMIP) is also very useful, yielding completely different results compared to classical MIP, and in particular, GMIP produces renderings that appear shaded (cf. images in Figures 3e & 3k). The *product* of both the classified gradient magnitude and the original signal can also yield better results. For instance, Figure 3i shows an enhancement of the engine metallic plate, which would be difficult to do just by summing the gradient magnitude (3h).

These results also show that using an affine transfer function in sum-based renderings, instead of a linear transfer function, can greatly enhance image contrast (Cf. spine visualization in Figure 4i). Therefore, this technique produces better results than regular Fourier volume rendering and is a better choice unless real-time performance is needed.

### 4.2 Stereo Visualization within OIVR

From both our experience and the experience of other users, we have come to several conclusions. First, stereo perception from rendering using the *maximum* compositing function is usually very good, while understanding object

depth from summation renderings depends more on the rendering conditions and parameters. The number of volume samples contributing to each pixel can explain this fact. In MIP, there is only one sample contributing to a pixel, and this sample has a greater chance to be on both stereo images, since it has a high intensity. Therefore, the depth of most objects in the *maximum*-based renderings of Figure 3 is easily understandable.

In X-ray like stereo renderings, many samples are going to contribute to the final pixel, usually resulting in cloudy objects. This is particularly visible when visualizing images like 3a and 3j, where depth is noticeable, but requires time for the eye to accommodate. If the number of samples contributing to the pixel is reduced, then the cloudy effect is less important, as in gradient magnitude based sum renderings. Especially, seeing rendering 3h in stereo illustrates very well the advantage of this approach: the depth relationships of the objects are immediately clear. Depth perception is also very good in sharp triangular TF based renderings. Using stereo in the sinus (4e) and colon (4h, 4i) renderings suppresses most depth ambiguities caused by overlapping structures.


Finally, we have used stereo MIP on a blood vessel dataset and noticed an interesting interpretation made by our visual system. Let's imagine two crossing blood vessels (cf. Figure 5), with the one in front having an intensity less

than the one in back. Therefore, the one in back will be visible at the intersection and will hide the front vessel. However, we observed that in this case, the front vessel seems to appear most of the time as a *transparent* object. This phenomenon is probably due to both some imperfection of the visual system enhancing lines and contours, and the fact that the vessel's missing part is partially visible on one of the two stereo pairs. This clearly demonstrates how our brain can subconsciously increase perception of non-realistic renderings.

These results show our enhancements to OIVR rendering can greatly increase the spatial understanding of volumetric datasets compared to traditional MIP and X-ray images. Stereo with OIVR is more effective in resolving depth relationships of overlapping, but non-occluding objects. In fact, stereo is even more important in OIVR rendering than in traditional realistic volume rendering because the latter provides enough depth cues for volumetric understanding, while traditional OIVR lacks all depth cues.

#### 4.3 Multi-channels Renderings and Non-Linear TFs

Figure 4 shows some renderings from combining four OIVR models using four non-linear transfer functions. We also included two FVR-like renderings of the original signal (4a) and the gradient magnitude (4c).



**Figure 4:** Four channels renderings from our system. In all cases, but 4k, the post matrix is the identity matrix, resulting in transfer functions (violet areas) that are representative of the final mapping of colors.

Summing the original signal on four different channels, with a different TF for each channel, generated rendering 4b. The blue portion is made from a sharp triangular TF, emphasizing the object's surface. Another TF emphasizes soft tissues (in red) while the two other TFs show bone structures. The advantage of having a TF that sums only the bone densities is clearly shown and the entire bone structure is clearly discernible.

Images 4e, 4f and 4g show renderings that only use the sharp triangular portion of the previous transfer function. The pulsation's location is close to zero, where no volume sample is present. Therefore, the (iso)surface in contact with air (0 density) is highlighted. An interesting fact is that for the first time, the UNC head tongue and sinuses are clearly shown in volume rendered images. It also demonstrates that some isosurfaces can be easily rendered just from the original signal and a *sum* operator. Finally, our local lighting model greatly reduces the confusing rendering of object overlap areas. Image 4f shows a tongue enhancement using a spherical lighting with both a linear attenuation and a cut-off, while image 4g replaces Euclidian distance attenuation by distance to plane attenuation. In both renderings, the tongue visualization is clearly improved.

Another useful application of this transfer function could be virtual endoscopy [HMKB97], where the whole colon structure can be quickly distinguished, as visible in renderings 4h and 4i (the spine was also enhanced with a red TF). Our renderings from the colon dataset show highly contrasted and featured images, but we do not have specialist feedback at this point. However, we expect that these visualizations would avoid the long colon navigation required when using isosurfaces and could augment traditional techniques as a quick preview for localizing exploration. Notice that these renderings were generated with a very simple TF that could be automatically computed from the volume histogram, increasing the usability of the system [BPS97].

Examples of renderings using four gradient magnitude sums with different TFs are shown in Figures 4d and 4j. In these cases, the different TFs increase the discrimination of the different material surfaces.

Finally, blending 3 different models with quasi-linear TFs generated the image in Figure 4k. The white objects are rendered from a MIP model, highlighting high values of the signal, while the orange surfaces are generated from a post-classified gradient sum. The signal sum was also used to emphasize filled regions of the volume (Notice the pink cloudy effect).

#### 4.4 Limitations and Possible Improvements

Our system is designed to produce high-quality images. Therefore, floating-point frame buffers (required for precision) and long fragment programs have been used, which decrease performance. We are working on optimizing our rendering performance and our fragment programs and also exploring the use of pre-integration as another possible performance solution [EKE01].

We are also exploring ways to improve our models through the use of two-dimensional transfer functions [KKH02] and improving the rendering of areas of excessive object overlap through the use of automatic attenuation.

## 5. Conclusion

Our research has demonstrated that with our new improvements, order independent volume rendering can be a very effective tool. While eliminating alpha blending allows better global understanding of the entire volume, the simple but non-previously mentioned use of stereo rendering can recover the depth information that was previously missing from such models. We have made several extensions to the traditional OIVR techniques (MIP, X-ray), including MIP of gradient magnitude and using sharp triangular transfer functions that allows shaded transparent isosurface visualization without requiring storage of additional information per voxel. Moreover, by handling up to four OIVR models in the pipeline, the system can improve the image perception, while keeping the familiarity within the medical community of MIP and X-ray type images.

This work could also be extended to illustration renderings, where the use of summed transparency and stereo for conveying information should be very useful. We plan to explore improvements to stereo OIVR techniques and conduct user studies of the effectiveness of the presented models in increasing depth perception in visualization.

## Acknowledgements


This material is based upon work supported by the US National Science Foundation under grants: NSF ACI-0081581, NSF ACI-0121288, NSF IIS-0098443, NSF ACI-9978032, and NSF ACI-0222675. We would like to thank Dirk Bartz for quickly providing us the virtual endoscopy dataset, Kirk Riley for helping us to set up OpenGL extension, and NVidia for providing us a Quadro FX 3000.

## References


- [BMDF02] Bergner S., Möller T., Drew M. S., Finlayson F. D.: Interactive spectral volume rendering. In *Proc. IEEE Visualization 2002*, pp. 101-108.
- [BPS97] Bajaj C. L., Pascucci V., Schikore D. R.: The contour spectrum. In *Proc. IEEE Visualization 1997*, pp. 167-173.
- [CCF94] CABRAL B., CAM N., FORAN J.: Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proc. of the 1994 symposium on Volume visualization*, pp. 91-98.
- [CMH\*01] CSÉBFALI B., MROZ L., HAUSER H., KÖNIG A., GRÖLLER E.: Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum* 20, 3 (Sept. 2001), C452-C460 (Proc. Eurographics 2001).

- [CSK03] CSÉBFALI B., SZIRMAI-KALOS L.: Monte Carlo Volume Rendering. In *Proc. IEEE Visualization 2003*, pp. 449-456.
- [EKE01] ENGEL K., KRAUS M., ERTL T.: High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *Proc. of Eurographic Workshop on Graphics Hardware'01*. pp. 9-16.
- [HBH03] HADWIGER M., BERGER C., HAUSER H.: High Quality two-level volume rendering on segmented data sets on consumer graphics hardware. In *Proc. IEEE Visualization 2003*, pp. 301-308.
- [HJC03] HÖNIGMANN D., RUISZ J., HAIDER, C.: Adaptive design of a global opacity transfer function for direct volume rendering of ultrasound data. In *Proc. IEEE Visualization 2003*. pp. 489-496.
- [HMBG01] HAUSER H., MROZ L., BISCHI G.-I., GRÖLLER E.: Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (2001), 242-252.
- [HMKB97] HONG L., MURAKI S., KAUFMAN A., BARTZ D., HE, T.: Virtual Voyage: Interactive Navigation in the Human Colon. In *Proc. Siggraph'97* (1997), pp 27-34.
- [KVH84] KAJIYA J., VON HERZEN B.: Raytracing Volume Densities. In *Proc. Siggraph'84* (1984), pp. 165-174.
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multi-Dimensional Transfer Functions for Interactive Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 8, 4 (July 2002), 270-285.
- [KPI\*03] KNISS J., PREMOZE S., IKITS M., LEFOHN A., HANSEN C., PRAUN P.: Gaussian Transfer Functions for Multi-Field Volume Visualization. In *Proc. IEEE Visualization 2003*, pp. 497-504.
- [LEV88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Application*. 9, 3. (1988), 29-37.
- [LEV90] LEVOY M.: Efficient raytracing of volume data. *ACM Transaction On Graphics* 9, 3, 1990, 245-261.
- [LM02] LUM E. B., MA K.-L.: Hardware-accelerated parallel nonphotorealistic volume rendering. In *Proc. of NPAR 2002*, pp. 67-74
- [LME\*02] LU A., MORRIS, C., EBERT D., RHEINGANS P., HANSEN C.: Non-photorealistic volume rendering using stippling techniques. In *Proc. IEEE Visualization 2002*, pp. 211-218.
- [MAL93] MALZBENDER T.: Fourier domain volume rendering. *ACM Transaction On Graphics* 12, 3, 1993, 233-250.
- [MAX95] MAX N.: Optical model for direct volume rendering. *IEEE Transaction On Visualization And Computer Graphics*, 1, 2 (1995), pp. 99-108.
- [MHG00] MROZ L., HAUSER H., GRÖLLER E.: Intensity high-quality maximum intensity projection. *Computer Graphics Forum* 19, 3 (Sept. 2000), C341-C350 (Proc. Eurographics 2000).
- [PSL\*95] POSTON T., SERRA L., LAWTON W., CHUA, B.C.: Interactive tube finding on a Virtual Workbench. In *Proc. Second International Symposium on Medical Robotics and Computer Assisted Surgery*, pp. 119-123.
- [RE01] RHEINGANS P., EBERT D.: Volume Illustration: Nonphotorealistic Rendering of Volume Models. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (2001), 253-264.
- [SC02] SHAREEF N., CRAWFIS R.: A view-dependent approach to MIP for very large data. In *Proc. SPIE EI '02* 4665 (2002), pp. 13-21.
- [SE03] SVAKHINE N., EBERT D.: Interactive volume illustration and feature halos. In *Proc. Pacific Graphics 2003*, 2003, pp. 347.354.
- [SSNK98] SATO Y., SHIRAGA N., NAKAJIMA S., TAMURA S., KIKINIS R.: LMIP: Local maximum intensity projection - a new rendering method for vascular visualization. *JOURNAL OF COMPUTER ASSISTED TOMOGRAPHY* 22, 6, 1998, 912-917.
- [TL93] TOTSUKA T., LEVOY M.: Frequency domain volume rendering. In *Proc. SIGGRAPH'93*, pp. 271-278.
- [WE98] WESTERMANN R., ERTL T.: Efficiently Using Graphics Hardware in Volume Rendering Applications. In *Proc. Siggraph'98* (1998), pp. 169-177.
- [WKE02] WEILER, M., KRAUS, M., ERTL, T.: Hardware-Based View-Independent Cell Projection. In *Proceedings Symposium on Volume Visualization '02*, 2002, pp. 13-22.





**Figure 5:** Virtual endoscopy dataset. The stereo image (left) requires red (left eye) and green (right eye) glasses.



**Figure 6:** Stereo Maximum Intensity Projection of a blood vessel.