

# Per-pixel Opacity Modulation for Feature Enhancement in Volume Rendering

Stéphane Marchesin and Jean-Michel Dischler and Catherine Mongenet  
LSIIT – Louis Pasteur University – Strasbourg, France

**Abstract**—Classical direct volume rendering techniques accumulate color and opacity contributions using the standard volume rendering equation approximated by alpha blending. However, such standard rendering techniques, often also aiming at visual realism, are not always adequate for efficient data exploration, especially when large opaque areas are present in a dataset, since such areas can occlude important features and make them invisible. On the other hand, the use of highly transparent transfer functions allows viewing all the features at once, but often makes these features barely visible. In order to enhance feature visibility, we present in this paper a straightforward rendering technique that consists of modifying the traditional volume rendering equation. Our approach does not require an opacity transfer function, and instead is based on a function quantifying the relative importance of each voxel in the final rendering called *relevance function*. This function is subsequently used to dynamically adjust the opacity of the contributions per-pixel. We conduct experiments with a number of possible relevance functions in order to show the influence of this parameter. As will be shown by our comparative study, our rendering method is much more suitable than standard volume rendering for interactive data exploration at a low extra cost. Thereby, our method avoids feature visibility restrictions without relying on a transfer function and yet maintains a visual similarity with standard volume rendering.

## I. INTRODUCTION

Traditional direct volume rendering techniques consist in integrating opacity and color values along a viewing ray by using a simple light transport model, called the volume rendering equation, inspired by the physics of light traversing a participating media (usually only absorption is considered, with no scattering). Using this approach, direct volume rendering allows one to display multiple values contributing to one pixel, unlike iso-surface extraction methods, which only consider single values. Therefore, volume rendering is widely used for the interactive exploration of 3D datasets since it allows one to display all data "at once". However, one major issue consists of defining a transfer function which associates opacity and color values to the data. Setting up this transfer function and in particular the opacity function is an intricate task, since it generally requires the expertise of the user and is deeply data dependent.

Setting the transfer function can be done either by hand, which is, however, an excessively time-consuming and painstaking process, or automatically or semi-automatically [8], [18], by introducing a specific data analysis process. In either case, it cannot be excluded that for some datasets wrong decisions could be made when building the transfer function. This has implications for the common volume rendering techniques:

- In the case of direct volume rendering, opaque features might occlude some other parts of the dataset for some or even,

in the worst case, for all viewing directions. This could make feature visibility in the dataset highly view-dependent. Conversely, setting up a transfer function with insufficient opacity might result in unintelligible pictures. Consequently, classical volume rendering techniques, especially based on the traditional volume rendering equation, might become impractical for efficient data exploration.

- In the case of the additive blending technique, the produced pictures are often under or over saturated. Since this technique uses a simple additive blending equation, over saturation of pictures can happen very quickly and with only a very limited number of contributions.

In order to resolve the crucial issue of simultaneous visibility of data features in a view-independent fashion, one can resort to a volume rendering integral modification in the hope of avoiding a potentially incorrect transfer function setup. By doing so, these methods are trading visual realism for pictures providing a better understanding of the data.

In a previous paper [11], we introduced such a method, that consists in modifying the volume rendering integral to achieve better visibility of the internal structures in datasets. It is view-dependent and dynamically adjusts the volume rendering integral per-pixel, thereby improving the resulting pictures by conveying more dataset features at once. In order to identify empty areas, a binary classification of the data was required. In this paper, we refine this principle by further introducing a function called *relevance function* that quantifies the relative importance of each voxel in the final rendering. We experiment with different suitable alternatives for this function, such as using the gradient and second order derivative. As the approach is still automatic, no complex transfer function setup is required. Unlike completely non-photo realistic volume rendering techniques, our approach remains very close to traditional volume rendering, thus providing pictures close to usual visual realism. Yet it can be considered a non-photo realistic method since it does not respect the actual physical light transport equation.

In the next section, we detail related works. In section 3, we present our new volume rendering technique. We then study the influence of different *relevance functions* in section 4. In section 5, we make some adjustments to improve visual quality. We discuss preintegration of our volume rendering equation in section 6. The implementation of our technique on the GPU is detailed in section 7 and we show it can reach interactive performance. We present results with various datasets and compare our technique to other volume visualization methods in section 8. Finally, we give concluding remarks in section 9.

## II. RELATED WORK

In recent years, many non-photorealistic volume visualization techniques have been proposed. These techniques trade picture

realism for greater data understanding.

Illustrative and non-photorealistic techniques have been used in the volume rendering field for more than a decade. In particular, numerous techniques allow prioritizing volume samples along a ray. Saito *et al.* [15] propose such a technique that achieves non-photorealistic previewing of volume fields using simple primitives such as points or lines which outline the shape of the volume structures. Illustrative techniques are also used by Ebert *et al.* [3] to enhance important features of the volume data. The authors propose local modifications of the volume rendering integral for feature enhancement and the addition of depth cues to increase depth perception in volume images. Viola *et al.* [19], [20] introduce a technique called importance-driven volume rendering. This technique uses pre-defined priorities between structures in the dataset and then renders accordingly: low priority structures can be occluded by higher priority ones, while high priority structures cannot be occluded by lower priority ones. However, this requires *à priori* knowledge of the dataset in order to segment and prioritize the object features. In order to keep context information available, Bruckner *et al.* [1] propose a model preserving contextual information while prioritizing relevant information according to real-time adjustable parameters. Thereby, it is possible to create visualizations depicting both important and contextual structures, where the latter will not hide the former. Csébfalvi *et al.* [2] present a visualization technique which uses gradient-based voxel selection, and only renders the relevant voxels. This results in efficient visualization of contours in the dataset. Similarly, Šereda *et al.* [17] propose a technique using the LH histogram allowing boundary visualization in volume datasets.

To achieve better visibility of volume structures, work has been undertaken to modify the volume rendering integral. Kraus [9] introduces scale-invariant volume rendering. This technique integrates the data in data space, and therefore achieves scale invariance for volume visualization. By its design, this technique leads to the same color and opacity for structures bearing the same density but different thicknesses. Sato *et al.* [16] modify the maximum intensity projection scheme by selecting the first value above a threshold on a given ray instead of the maximum value over the whole ray. Hauser *et al.* [7] mix two well-known volume visualization techniques, namely direct volume rendering and maximum intensity projection, into a joint method. The dataset is first segmented into two classes, and those classes are then rendered using either one or the other technique. Order independent volume rendering techniques are explored by Mora *et al.* [12]. Such techniques extend existing maximum intensity projection (MIP) [13], [14] and X-ray techniques, and do not require sorting of the volume primitives. In order to improve the visibility of boundaries within the volume, Lum *et al.* [10] propose lighting based transfer functions.

In order to improve the quality of numerical integration, Engel *et al.* [4] propose a technique called preintegration. This technique increases the visual quality of volume rendering by pre-computing slabs of the volume rendering integral.

In the field of high dynamic range (HDR) volume visualization, Ghosh *et al.* [5] have developed a technique to achieve volume rendering on high dynamic range displays. Yuan *et al.* [21] visualize high dynamic range datasets on a standard display system using two passes: first the data is rendered into a high dynamic range buffer, and then the resulting image is tone mapped into an image suitable for display.

In the field of 2D imaging, the local histogram equalization technique [6] enhances the contrast of a picture by performing histogram equalization in the neighbourhood of each pixel.

Even though locally adaptive techniques are widely used in the image processing field, such techniques have been rarely used for visualization. Among all the techniques discussed above, very few take advantage of view-dependent rendering to increase the amount of information in the picture. Those that do so usually require the user to go through a complex data segmentation and/or transfer function setup phase. Since this setup is view-independent, it does not adapt to the current viewing conditions in order to maximize visibility of internal structures. Furthermore, the complexity of the preprocessing phase usually makes volume rendering unsuitable as a data exploration tool. In this paper, our purpose is to move that complex phase further in the visualization pipeline by dynamically and locally adjusting the rendering at runtime. We also aim at dataset exploration, and thus our algorithm does not require any complex setup, and is able to achieve interactive rendering. Finally, as opposed to Kraus [9], we want to be able to distinguish objects with different thicknesses and therefore thin objects should not appear similarly to thick ones. Therefore, we introduce a new approach which is halfway between classical photo-realistic volume visualization techniques and non-photorealistic volume visualization techniques. It locally adjusts the opacity values of the contributions, and thereby manages to keep a greater number of features visible at the same time. Since opacity values are computed on the fly, our method does not need any opacity transfer function setup, but only a binary thresholding of empty areas.

### III. REFORMULATING VOLUME RENDERING USING A RELEVANCE FUNCTION

The volume rendering technique we present in this section aims at solving the visibility issues pertaining to volume rendering by dynamically adapting the opacity of the fragment contributions at a per pixel level. Our technique is based on the use of a *relevance function* quantifying the relative importance of each voxel in the final rendering according to a given number of parameters. This function is subsequently used to compute the voxel opacity, and therefore removes the need for an opacity transfer function setup, which can prove time-consuming for the user. Such a function depends only on the data, and can be realized in various ways, based on intrinsic signal properties like the scalar value, or the local gradient. Let us now present our per pixel adaptive technique in both contexts of direct volume rendering (DVR) and additive volume rendering.

#### A. Per-pixel adaptive DVR

Let  $C(x, y)$  represent the final pixel color at  $(x, y)$ ,  $s(i)$  the scalar value of the  $i^{\text{th}}$  sample along a viewing ray cast from  $(x, y)$ ,  $c()$  the luminance transfer function for scalar  $s$  so that  $c(s(i))$  is the luminance of the  $i^{\text{th}}$  sample,  $\tau()$  the extinction function so that  $\tau(s(i))$  is the extinction of the  $i^{\text{th}}$  sample, and  $L$  the length of the ray. The classical direct volume rendering integral is then as follows:

$$C(x, y) = \int_0^L c(s(t))\tau(s(t))e^{-\int_0^t \tau(s(u))du} dt \quad (1)$$

It can be approximated by the following Riemann sum:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} \tilde{c}(\tilde{s}(i)) \tilde{\tau}(\tilde{s}(i)) \prod_{j=0}^{i-1} (1 - \tilde{\tau}(\tilde{s}(j))) \quad (2)$$

Let us now define the *relevance* of a contribution: we define the *relevance* of a contribution as its importance relative to other contributions. Figure 1 shows the simple case of a binary *relevance* function where *relevant* contributions are samples that fall within segmented parts of the datasets (shown in blue on the figure). Let  $\delta(\cdot)$  be our classification function. Intuitively, this




Fig. 1. Counting the relevant contributions (in bold) carried by a ray going through a segmented dataset (in blue). Among the 22 samples, there are only 17 which are relevant.

function classifies the *relevance* of a given contribution over a ray. For the beginning let us assume that this function takes a binary value defined inside the volume, so that  $\delta(i)$  is 1 if the  $i^{\text{th}}$  sample is relevant and 0 otherwise. In order to keep all structures visible that project to a given pixel, we use the same extinction coefficient  $\tilde{\tau}_0$  for each term. Then the product in Equation 2 becomes a power with exponent  $\sum_{j=0}^{i-1} \delta(j)$ . If we furthermore multiply each contribution with the relevance weight, we get:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} \delta(i) \tilde{c}(\tilde{s}(i)) \tilde{\tau}_0 (1 - \tilde{\tau}_0)^{\sum_{j=0}^{i-1} \delta(j)} \quad (3)$$

In order to keep as many features as possible visible at once, we compute the value of  $\tilde{\tau}_0$  that maximizes the visibility of the furthest contribution on the screen. Let us denote the sum of the relevance coefficients as  $\Delta = \sum_{j=0}^{n-1} \delta(j)$ . The total extinction for the furthest contribution is:

$$\tilde{\tau}_{\text{furthest}}(x,y) = \tilde{\tau}_0 (1 - \tilde{\tau}_0)^\Delta \quad (4)$$

To maximize this function, we take its derivative and find its zeroes:

$$\frac{d\tilde{\tau}_{\text{furthest}}(x,y)}{d\tilde{\tau}_0} = (1 - \tilde{\tau}_0)^\Delta - \frac{\tilde{\tau}_0 (1 - \tilde{\tau}_0)^{\Delta-1} \Delta}{1 - \tilde{\tau}_0} \quad (5)$$

We are only interested in extinction values in  $]0, 1[$  since  $\tilde{\tau}_0 = 0$  and  $\tilde{\tau}_0 = 1$  respectively mean fully transparent and fully opaque:

$$(1 - \tilde{\tau}_0)^\Delta - \frac{\tilde{\tau}_0 (1 - \tilde{\tau}_0)^{\Delta-1} \Delta}{1 - \tilde{\tau}_0} = 0 \quad (6)$$

$$\Rightarrow (1 - \tilde{\tau}_0)^{(\Delta+1)} - \tilde{\tau}_0 (1 - \tilde{\tau}_0)^{\Delta-1} \Delta = 0 \quad (7)$$

$$\Rightarrow (1 - \tilde{\tau}_0)^{\Delta+1} = \tilde{\tau}_0 (1 - \tilde{\tau}_0)^{\Delta-1} \Delta \quad (8)$$

$$\Rightarrow (1 - \tilde{\tau}_0) = \tilde{\tau}_0 \Delta \quad (9)$$

$$\Rightarrow \tilde{\tau}_0 = \frac{1}{\Delta + 1} \quad (10)$$

In order to maximize visibility for the furthest contributions, one should therefore choose for  $\tilde{\tau}_0$  the value  $\frac{1}{\Delta}$ .

Now, substituting in Equation 2 the variable  $\tau(\tilde{s}(i))$  with  $\tilde{\tau}_0 \times \delta = \frac{\tilde{\tau}_0}{\Delta+1}$ , we finally end up with the following rendering formula:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} c(\tilde{s}(i)) \frac{\delta(i)}{\Delta+1} \prod_{j=0}^{i-1} \left(1 - \frac{\delta(j)}{\Delta+1}\right) \quad (11)$$

Note that there is no longer any  $\tau(\cdot)$  extinction function in this formula.

This formula has been computed for a binary relevance function  $\delta$ . We may generalize this expression to continuous  $\delta(i)$  functions with values in  $[0, 1]$ , so that  $\delta(i)$  gets closer to 1 as the  $i^{\text{th}}$  sample becomes more *relevant*, and closer to 0 when it becomes less *relevant*. To do that, and for the sake of simplicity, we propose to use the same Equation 11 as an approximation of what an explicit substitution of  $\tau(\cdot)$  by  $\tilde{\tau}_0 \times \delta$  would give in this case.

### B. Per-pixel additive volume rendering

Let  $C(x,y)$ ,  $c(\cdot)$ ,  $\delta(\cdot)$  and  $\tau(\cdot)$  be defined as previously. The additive blending volume rendering integral is as follows:

$$C(x,y) = \int_0^L c(s(t)) \tau(s(t)) dt \quad (12)$$

which can be approximated by a Riemann sum in the following way:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} c(\tilde{s}(i)) \tilde{\tau}(\tilde{s}(i)) \quad (13)$$

As before, we multiply each term by  $\delta(i)$  to introduce classification in the context of additive volume rendering:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} \delta(i) c(\tilde{s}(i)) \tilde{\tau}(\tilde{s}(i)) \quad (14)$$

In order to avoid saturating colors at a given pixel with our additive technique, we use the same  $\tau$  for all contributions, in a way that averages the values of the *relevant* contributions gathered over a ray according to their respective *relevance* values:

$$\tau_0 = \frac{1}{\Delta} \quad (15)$$

We now replace  $\tau(\cdot)$  with  $\tau_0$  in Equation 14 and obtain:

$$\tilde{C}(x,y) = \frac{1}{\Delta} \sum_{i=0}^{n-1} c(\tilde{s}(i)) \delta(i) \quad (16)$$

which can be rewritten as:

$$\tilde{C}(x,y) = \sum_{i=0}^{n-1} \frac{c(\tilde{s}(i)) \delta(i)}{\Delta} \quad (17)$$

A continuous formulation of this equation is given by:

$$C(x,y) = \frac{\int_0^L c(s(t)) \delta(s(t)) dt}{\int_0^L \delta(s(t)) dt} \quad (18)$$

We will use this continuous formulation in the section concerning pre-integration.

Note that Equation 17 is invariant under scaling  $\delta(i) \rightarrow \alpha \delta(i)$  with some nonzero constant  $\alpha$ . For Equation 11, this is approximately true for  $\Delta \gg 1$ . In the following, we therefore drop the requirement of normalizing the relevance functions to the range  $[0, 1]$ .



#### IV. RELEVANCE FUNCTIONS

In this section, we study the influence of the relevance function  $\delta()$  on the final results, and we experiment with multiple relevance functions using different parameters.

- The simplest choice for  $\delta()$  consists of using a binary function. That is,  $\delta()$  is such that  $\delta(i) = 1$  if the contribution of the  $i^{\text{th}}$  sample is *relevant* and 0 otherwise according to the segmentation depicted on figure 1. In practice, we simply use a binary classification of the scalar value at each sample  $s(i)$ . This is a solution we have used in a previous paper [11].
- Another solution that we experimented with consists of weighting the relevant contributions instead of considering them as all being equal. A common assumption is to consider regions of high variance, e.g. high gradient, as more important than regions with low variance, e.g. low gradient. It therefore makes sense to use the gradient value as a relevance function. In this case, if  $f(x, y, z)$  is the function we want to visualize, the relevance of a contribution is defined as  $\delta(i) = \|\nabla f(x, y, z)\|$ .
- We have also tried to use a thresholding of gradient values, that is, the relevance function is defined as  $\delta(i) = \begin{cases} 1 & \text{if } \|\nabla f(x, y, z)\| > t \\ 0 & \text{otherwise} \end{cases}$  where  $t$  is a user-defined threshold.
- In a similar way to the first order gradient, we have conducted experiments using the second order derivative or Laplacian as a relevance criterion, that is  $\delta(i) = \Delta f(x, y, z) = \nabla \cdot \nabla f(x, y, z)$ .
- Feature boundaries are often considered important in datasets. These can also be taken into account by our relevance function by considering both the gradient and second order derivative. Contributions have an increasing gradient as they approach boundaries, while the Laplacian crosses zero near a boundary. We have implemented a simple boundary detection as follows:  $\delta(i) = \begin{cases} 1 & \text{if } \|\nabla f(x, y, z)\| > t_1 \text{ and } \|\Delta f(x, y, z)\| < t_2 \\ 0 & \text{otherwise} \end{cases}$ , where  $t_1$  and  $t_2$  are user-defined thresholds.

Some of these techniques for choosing relevance functions are fully automatic and require no user specific manipulation, some others require minimal intervention like setting a threshold up or defining a binary classification. We show various examples of these different functions in section 8.

#### V. METHOD ADJUSTMENTS

If the per-pixel volume rendering technique is used strictly as described previously, it produces images that do not clearly depict the object structures as shown in the second image of the top row of Figure 2. In order to efficiently use this method, one has to make some adjustments to it. The first of these adjustments bounds the extinction values. In order to achieve visual continuity, the second adjustment filters what we call the *relevance map*: the sum of *relevance* values for each pixel. If the relevance function  $\delta(i)$  is binary, it represents a simple count of the number of relevant contributions. Another adjustment is the addition of depth cues through depth-based coloring and the last one consists of adding a local shading in order to improve the perception of shapes. These last two adjustments produce pictures which have some visual relationship with traditional volume rendering. We now describe these adjustments in detail:

- When the extinction values are used as-is, our technique adjusts even the lowest extinction coefficients to higher values, which gives false visual clues. It occurs for instance when there are only few *relevant* contributions to a pixel, that is  $\sum_{i=0}^{n-1} \delta(\tilde{s}(i))$  is small and inversely  $\tilde{\tau}_0$  becomes high. We therefore decided to put an upper bound  $\tilde{\tau}_{max}$  to  $\tilde{\tau}_0$  by using  $\min(\tilde{\tau}_{max}, \tilde{\tau}_0)$  instead of  $\tilde{\tau}_0$ . We conducted a number of experiments with different datasets and found that using a maximum value prevents all areas from becoming opaque, and therefore improves depth perception inside the data. This bound was determined experimentally to be within the  $[0.1, 0.2]$  range. The third and fourth images of the top row of Figure 2 show  $\tilde{\tau}_{max} = 0.25$  and  $\tilde{\tau}_{max} = 0.12$ , respectively.
- In order to avoid giving false information about the dataset, spatial continuity should be ensured in screen space. However, when too few contributions are made to a single pixel, it can result in discontinuities in the picture which look like edges. To avoid this, we filter the relevance map using a Gaussian blur. After rendering the relevance map, we run it through a blurring filter that removes most of the high-frequency data. The second pass is then done from that same map. Thanks to that improvement, the noise from the relevance map is successfully removed and internal structures become clear as shown on the bottom left of Figure 2. An alternate implementation of filtering is to render the relevance map at a lower resolution during the first pass, and stretch it using the card's native bilinear filtering capabilities. Using this functionality is faster, at the expense of less accurate results as shown by the second and third images of the bottom row of Figure 2 which show 2 times and 4 times scaling, respectively. These pictures show that it is possible to greatly improve the interactivity of our technique at the expense of visual quality.
- In the case of additive volume rendering, the depth information is lost since the blending method is commutative. This means that two contributions, one on the back of the dataset, and the other on the front, will have the same result on screen. Therefore, it is primordial to restore the depth information. In this purpose, we add depth-based coloring of the fragment samples: as the samples get further from the observer, we modulate their color proportionally to the depth of the sample, i.e. we multiply the resulting color by a factor decreasing linearly with the depth of the sample. This allows an increased depth perception in the produced pictures as shown by the bottom right image of Figure 2.
- In order to improve the perception of shapes it is also important to add a local shading model, such as the Phong one. This consists of replacing the colour  $\tilde{c}(\tilde{s}(i))$  with the following formula:

$$\tilde{c}_l = \tilde{c}(\tilde{s}(i)) * (K_a + K_d \times \vec{L} \cdot \vec{N} + K_s \times (\vec{R} \cdot \vec{V})^\alpha)$$

where  $K_a, K_d$  and  $K_s$  are the Phong coefficients,  $\alpha$  is the specular exponent,  $\vec{L}$  is the light vector, the normal  $\vec{N}$  is assimilated to the gradient of  $\tilde{s}(i)$ ,  $\vec{V}$  is the direction towards the viewer and  $\vec{R}$  the local reflection vector computed according to the local gradient.

#### VI. PREINTEGRATION

Preintegration [4] is an important feature for volume rendering. It is possible to achieve preintegration in our context of per-






Fig. 2. (a) Relevance map (representing the sum of *relevant* contributions for each pixel), Pictures produced from: (b) naive implementation with  $\bar{\tau}_{max} = 1$  unfiltered (20.5 fps), (c)  $\bar{\tau}_{max} = 0.25$  unfiltered (20.5 fps), (d)  $\bar{\tau}_{max} = 0.12$  unfiltered (20.5 fps), (e)  $\bar{\tau}_{max} = 0.12$  Gaussian filtered (8.5 fps), (f)  $\bar{\tau}_{max} = 0.12$  bilinear with 2 times (24.5 fps) bilinear scaling, (g) 4 times (25.3 fps) bilinear scaling and (h) z-based coloring with  $\bar{\tau}_{max} = 0.12$  and Gaussian filtering (8.5 fps).

pixel additive volume rendering, provided the relevance function depends on a single scalar parameter. In this section, we describe how this can be done in the case of a relevance function depending on the scalar value. Starting from the additive adaptive volume rendering integral given by equation 18, we have:

$$C(x, y) = \frac{\int_0^L c(s(t))\delta(s(t))dt}{\int_0^L \delta(s(t))dt} = \frac{\sum_{i=0}^{n-1} \bar{c}(i)}{\sum_{i=0}^{n-1} \bar{\delta}(i)} \quad (19)$$

with

$$\bar{c}(i) = \int_{\frac{L}{n}i}^{\frac{L}{n}(i+1)} c(s(t))\delta(s(t))dt \quad (20)$$

and

$$\bar{\delta}(i) = \int_{\frac{L}{n}i}^{\frac{L}{n}(i+1)} \delta(s(t))dt \quad (21)$$

which can be approximated as follows by assuming a linear scalar function  $s()$  over the  $i^{th}$  interval:

$$\begin{aligned} \bar{c}(i) &\approx \int_{\frac{L}{n}i}^{\frac{L}{n}(i+1)} c(s(i)(1 - (t - \frac{L}{n}i)) + s(i+1)(t - \frac{L}{n}i)) \\ &\quad \tau(s(i)(1 - (t - \frac{L}{n}i)) + s(i+1)(t - \frac{L}{n}i)) \\ &\quad \delta(s(i)(1 - (t - \frac{L}{n}i)) + s(i+1)(t - \frac{L}{n}i))dt \end{aligned} \quad (22)$$

and

$$\bar{\delta}(i) \approx \int_{\frac{L}{n}i}^{\frac{L}{n}(i+1)} \delta(s(i)(1 - (t - \frac{L}{n}i)) + s(i+1)(t - \frac{L}{n}i))dt \quad (23)$$

It is therefore possible to preintegrate the values of  $\bar{c}(i)$  and  $\bar{\delta}(i)$  in two separate tables and thus achieve preintegration in the context of adaptive additive volume rendering. Figure 3 shows that doing so greatly improves the quality of the pictures.

Note however that it is not convenient to preintegrate adaptive direct volume rendering since the table would have 3 entries: the two scalar values as with standard preintegration, and the current extinction coefficient. This, in turn, would require a 3D texture to store the preintegration table which implies a high video memory usage. Preintegration of relevance functions using multiple scalar parameters further increases the dimensionality requirement for the tables. For example a table preintegrating a relevance function that depends on both the gradient and the Laplacian would require a 5 dimensions in the context of adaptive direct volume rendering.

## VII. IMPLEMENTATION

We have implemented per-pixel opacity modulation for volume rendering on graphics hardware for voxel datasets. Our hardware implementation is essentially a 2-pass raycasting-based approach using 3D textures. It is depicted in figure 4 and works as follows:

- Initially, a criterion is chosen that allows discriminating relevant voxels.
- In a **first pass**, the sum of the *relevance* values for each screen pixel is computed using the previously chosen criterion, thereby creating the relevance map. In order to sum the relevance values at each pixel, this pass does a simple GPU-based raycasting of the dataset into a frame buffer object. As the accumulation is done within the fragment shader, full accuracy is retained until the final storing phase.
- In a **second pass**, given the relevance map and the color transfer function, the final rendering is produced. The frame buffer object of the *relevance map* is bound to a texture and the dataset is rendered again as shown in Figure 5. For each pixel, the fragment program first uses the fragment coordinate information to retrieve the sum of the relevance values. Then, a GPU-based raycasting is done through the dataset. Again, full accuracy is retained throughout the




Fig. 3. Our per-pixel modulated volume rendering technique without preintegration (top) and with preintegration (bottom) using the same sampling rate.

raycasting thanks to the use of internal floating point register as accumulators.

### VIII. RESULTS

We have experimented per-pixel modulated volume rendering with both DVR and additive volume rendering, using multiple relevance functions. In all tests, the binary function used is based on the opacity function of the DVR rendering, and is 1 when the opacity function is greater than zero, and 0 otherwise. All thresholds have been set at 0.5.

Figure 6 shows pictures obtained using our technique with different datasets. The first column uses the  $256 \times 256 \times 128$  bonsai dataset, the second and third ones use the  $256 \times 256 \times 225$  head dataset, and the fourth one the  $128 \times 128 \times 128$  bucky ball dataset. Standard rendering is compared to our adaptive technique, and multiple relevance functions are tested. For all these pictures, a local lighting model was used, and a hand-tuned opacity function was created for classical DVR and additive rendering. As can be seen from these pictures, per-pixel modulated volume rendering can significantly enhance details. In particular, borders are made sharper when using a binary classification and DVR, while the other features remain visible. The relevance functions allowing accurate border detection (namely the gradient value and border detection functions), although simpler from an end user point of view, allow visibility of most of the internal structures, especially

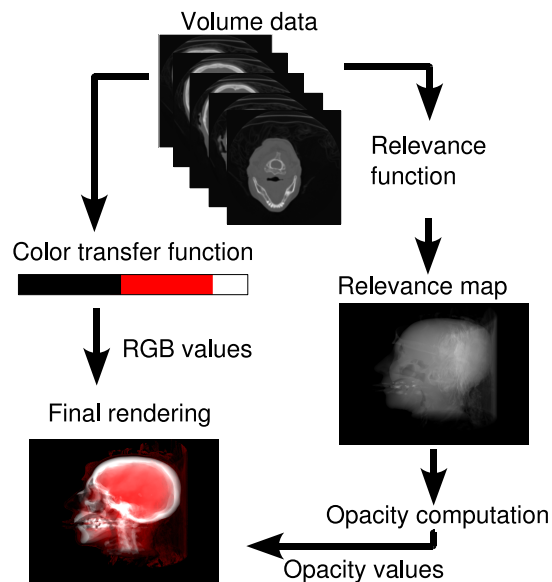


Fig. 4. Outline of our implementation.

when combined with additive rendering. For example, the use of the first gradient leads to pictures where most of the bonsai leaves are sharp. On the other hand, the use of the second gradient alone is not sufficient to correctly determine areas of interest inside the dataset, and leads to missing most of the structures. This is exemplified by the pictures from the bonsai dataset where not all leaves are visible, or from the bucky ball dataset where this criterion fails to detect most of the borders. Similarly, using a binary classification as the relevance function requires additional work, but also allows selecting what parts of the data should be visible (for example in the case of the head dataset, only the skin and bones were kept visible, so that others structures do not interfere with data understanding). In the case of additive volume rendering, per-pixel modulated volume rendering also globally avoids saturating pictures, but again keeps most details visible. This is shown on the head dataset renderings, where the borders of the bones and skin are more distinguishable than with the other approaches. Notice that the bucky ball dataset is synthetic, and does not feature clear boundaries between the different densities. Nevertheless, our technique is able to achieve good visibility of the internal structure of this dataset. One can also notice from figure 6 that our adaptive method performs similarly or in some cases better than classical volume rendering with a carefully hand-tuned opacity function, while at the same time it requires less user intervention.

Figure 7 compares per-pixel modulated rendering with the classical DVR technique using two opacity transfer functions and different viewpoints, for the  $379 \times 229 \times 305$  knee dataset. DVR is used with two opacity functions which are shown on top of the figure: a naive opacity function (the opacity transfer function is 0 everywhere except in the range  $[5, 85]$  where it is 1) and a hand-tuned opacity function that segments the skin and bones. The per-pixel modulated technique uses the naive opacity function as a binary classification. These pictures show that even though it uses a naive transfer function, our technique is able to convey most of the bone structure, similarly to what is obtained with a hand-tuned transfer function and DVR. As can be seen on the left column, DVR with a naive transfer function does not allow clearly

```

uniform sampler2D transfer_function;
uniform sampler2D relevance_map;
uniform sampler3D dataset;
uniform vec3 observer;
uniform vec3 sampling_distance;
void main()
{
    vec3 pos=gl_TexCoord[0].xyz;
    float scalar1;
    float scalar2;
    vec4 result = vec4(0.0, 0.0, 0.0, 1.0);
    vec3 progression;
    vec3 steps1 , steps2 , steps;
    vec4 frag_color;
    vec4 data_sample;
    int nr_steps;
    int i , j;
    float xc;
    float yc;
    float sum_relevance;
    float relevance;

    // compute the progression vector inside the data
    progression.xyz=gl_TexCoord[1].xyz-observer.xyz;
    progression.xyz*=(sampling_distance.xyz/length(progression.xyz));

    // compute the number of steps inside the data
    steps1.xyz=abs((1.0-pos.xyz)/progression.xyz);
    steps2.xyz=abs((pos.xyz)/progression.xyz);
    if (progression.x>0.0)
        steps.x=steps1.x;
    else
        steps.x=steps2.x;
    if (progression.y>0.0)
        steps.y=steps1.y;
    else
        steps.y=steps2.y;
    if (progression.z>0.0)
        steps.z=steps1.z;
    else
        steps.z=steps2.z;

    nr_steps=int(min(steps.x,min(steps.y,steps.z)));

    xc=gl_FragCoord.x*%f;
    yc=gl_FragCoord.y*%f;
    // look up the sum of the relevance function at the current pixel
    sum_relevance=(float(texture2D(relevance_map,vec2(xc,yc)));

    data_sample=texture3D(dataset, pos.xyz);
    scalar2=data_sample.a;
    pos.xyz+=progression.xyz;

    // raycast through the data
    for(j=0;j<256;j++)
    {
        for(i=0;i<256;i++)
        {
            scalar1=scalar2;
            data_sample=texture3D(dataset, pos.xyz);
            scalar2=data_sample.a;
            frag_color.rgba=texture2D(transfer_function,vec2(scalar2,scalar1));
            // find the relevance number for this fragment
            relevance=frag_color.a;
            // modulate the opacity of the fragment
            frag_color.rgba*=relevance/sum_relevance;

            // accumulate using additive blending
            result.rgba+=frag_color.rgba*vec4(relevance*sampling_distance);

            // step into the data
            pos.xyz+=progression.xyz;

            if (j*256+i>nr_steps) break;
        }
        if (j*256+i>nr_steps) break;
    }

    gl_FragColor.rgba=result.rgba;
}

```

Fig. 5. Shader code for the second pass of per-pixel modulated additive volume rendering. A simple binary classification criterion is used.

distinguishing the bone structure, whereas the adaptive approach with this same naive function makes this structure visible, in a way quite similar to the case of a hand-tuned transfer function. Furthermore, figure 7 also shows that viewpoint changes do not impact visibility of features with our technique, as opposed to DVR: for example the bones remain visible from all viewpoints.

Figure 8 compares per-pixel modulated rendering with maximum intensity projection (MIP) and a gradient-based relevance function using three datasets. Since maximum intensity projection is another explorative technique that does not require transfer function setup, it makes sense to include it in this comparison. These pictures show that per-pixel modulated volume rendering reveals more details from the dataset than the MIP technique. In

particular, the leaves of the bonsai and the structure of the bucky ball dataset are more easy to deduce from the adaptive pictures than from the MIP ones.

Figure 2 shows the influence of applying a  $3 \times 3$  Gaussian filter on the relevance map: small features that look like edges because of the adaptive opacity are successfully removed using this filter, while the rest of the features is still visible. Table I shows the performance of our technique compared to the classical volume rendering techniques. These measurements were conducted on a Pentium D 3.4GHz machine with a GeForce 7950 GT graphics card, and no preintegration was used. The classical DVR, additive and all the adaptive methods were all realized using GPU-based raycasting. Note that both DVR and additive volume rendering lead to the same performance results, and changing the relevance function results in a measured performance variation of less than 10%. These measurements show that our technique impacts the rendering interactivity only slightly, and maintain interactivity. These results also demonstrate that replacing the filtering stage by the combination of undersampling and bilinear filtering gives a significant performance gain. When bilinear upscaling by a factor of two is used, the relevance map to be computed is only one quarter of the relevance map that is required for Gaussian filtering. That, plus the additional filtering stage accounts for the performance difference.

	Classical	Bilinear 2x	Filter
Head	10.0 fps	7.4 fps	3.1 fps
Bonsai	15.8 fps	11.2 fps	4.6 fps
Bucky	23.2 fps	17.1 fps	7.0 fps
Knee	6.8 fps	5.2 fps	2.6 fps

TABLE I  
PERFORMANCE OF THE CLASSICAL VS PER-PIXEL MODULATED VOLUME RENDERING TECHNIQUES (NOT PREINTEGRATED).

## IX. CONCLUSIONS

In the context of data exploration, non-photorealistic techniques have shown that it is possible to increase the quality of the visualization by showing more data features. In this paper, we have introduced a new simple volume rendering technique that manipulates the opacity values in a view-dependent fashion in order to ensure maximal visibility of the internal data structures. We have compared this technique to other widely used volume rendering methods and hand-tuned transfer functions, and have demonstrated its efficiency. Our technique results in better understanding of the objects features, and furthermore does not require any complex opacity function setup. For some relevance functions it does not require any user intervention, while others require minimal user setup (either a binary classification or a single threshold value). We have found that using standard border detection techniques (gradient or border detection) as the relevance function leads to the best results, and without user input. Our method also ensures good visibility of the data features independently of the viewpoint, as opposed to the classical DVR method. Moreover, since our technique has been fully implemented on graphics hardware, we achieve interactive performance, thereby making it efficient in the context of data exploration, and allowing the user to use motion and interaction in order to better understand the internal structures of the dataset.



However, we think a lot of extensions are possible. First, we would like to extend the idea of per-pixel opacity to other volume rendering algorithms, in particular when multiple techniques are used at the same time (for example, DVR and isosurfaces). Second, we would like to derive more complex opacity modification functions, for example taking the depth position of the sample into account. Finally, we would like to experiment combining our technique with an automatic segmentation technique in order to form a fully automatic volume exploration tool.

#### X. ACKNOWLEDGEMENT

This work was supported by the French National Agency of Research (ANR) under project number ANR-05-MMSA-0008-01.

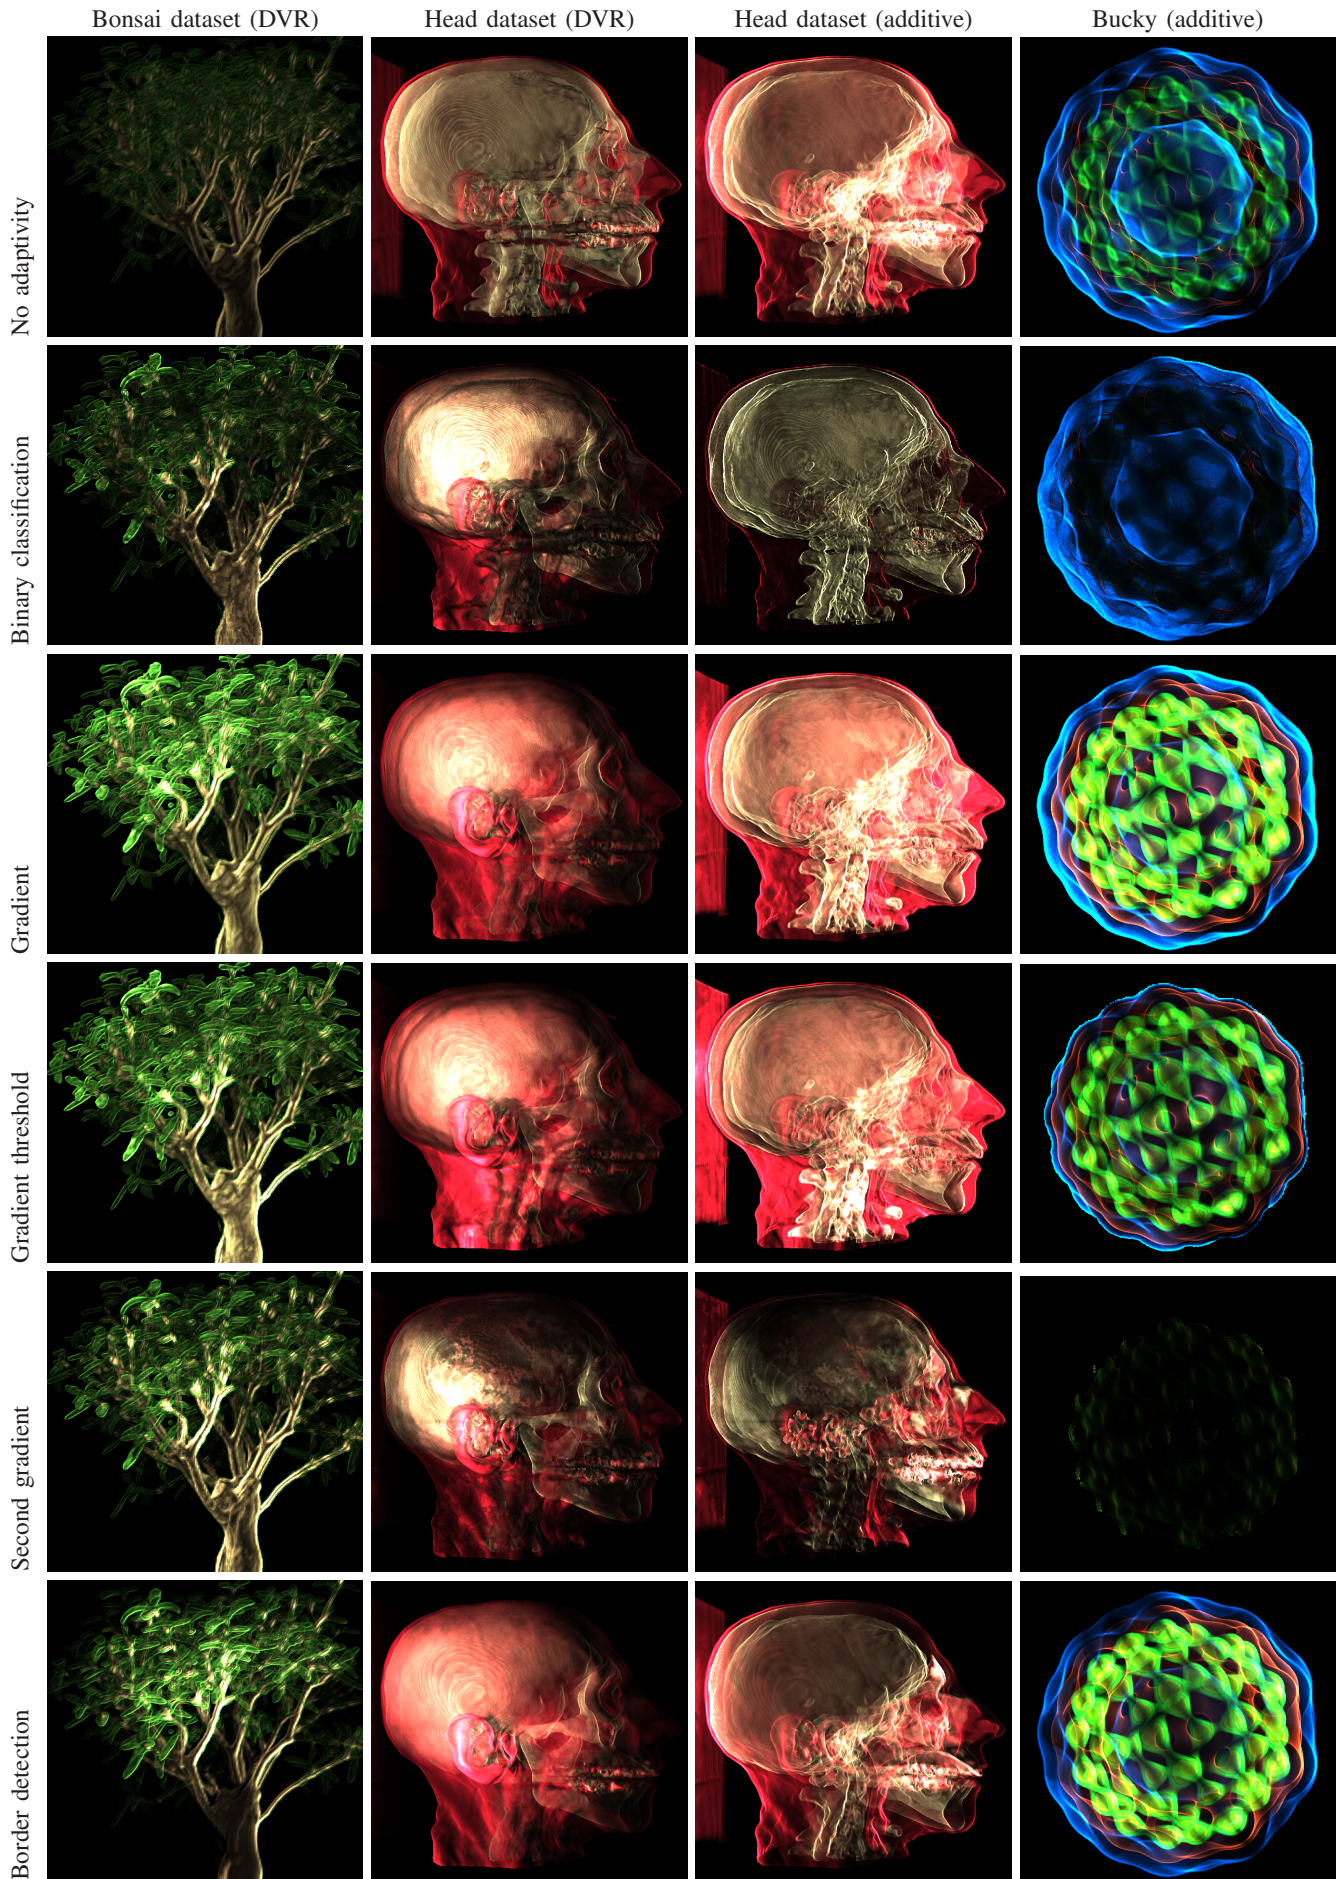


Fig. 6. Comparison of the additive and adaptive volume rendering techniques using different relevance functions.



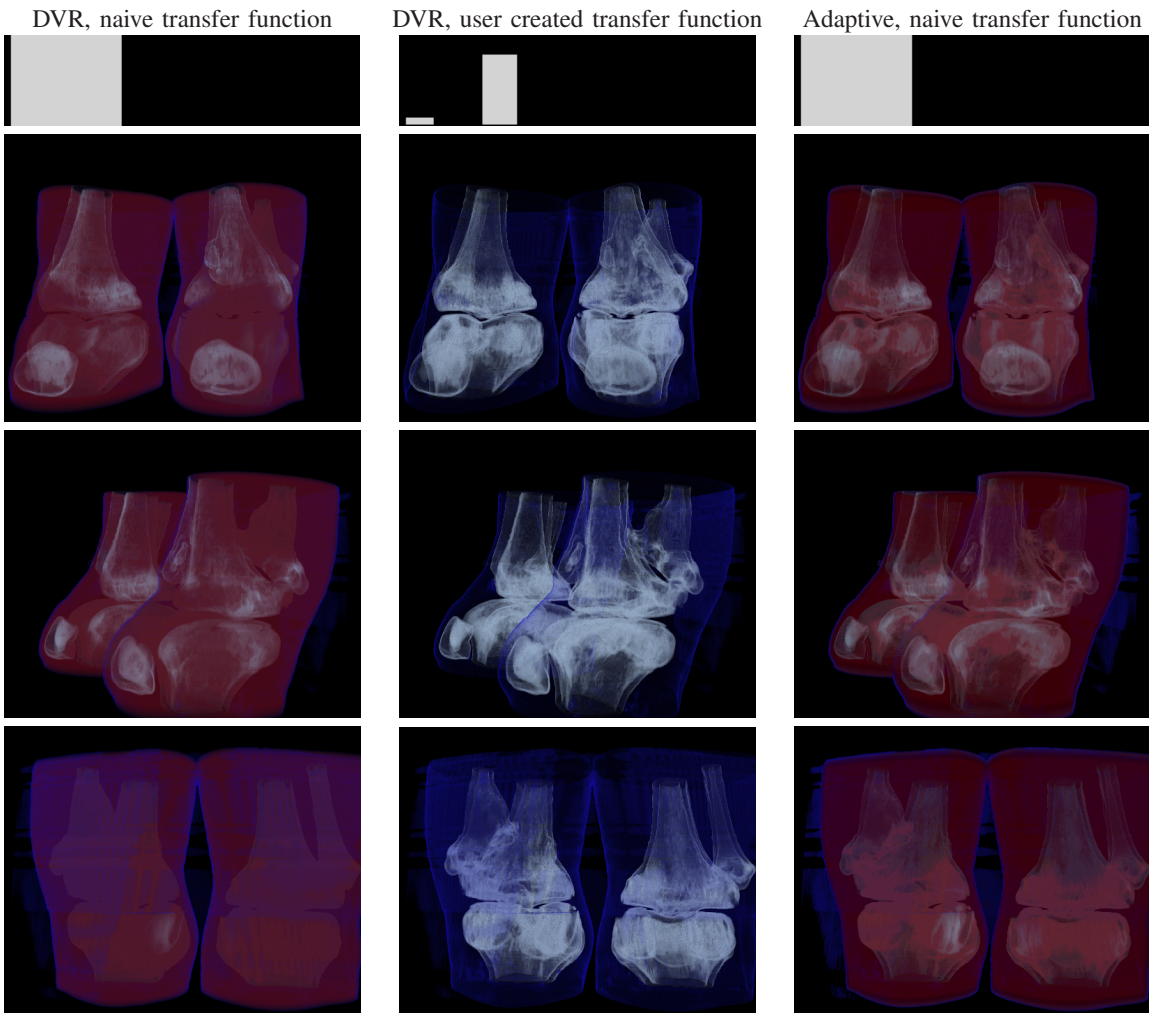


Fig. 7. Comparison of DVR (naive and hand-tuned transfer functions) with per-pixel modulated volume rendering (naive binary classification).




Fig. 8. Comparison of the per-pixel modulated volume rendering with maximum intensity projection (MIP)



## REFERENCES

- [1] Stefan Bruckner, Sören Grimm, Armin Kanitsar, and Meister Eduard Gröller. Illustrative context-preserving volume rendering. In *Proceedings of EuroVis*, pages 69–76, 2005.
- [2] Balázs Csebfalvi, Lukas Mroz, Helwig Hauser, Andreas König, and Meister Eduard Gröller. Fast visualization of object contours by non-photorealistic volume rendering. In *Eurographics*, pages 452–460, 2001.
- [3] David Ebert and Penny Rheingans. Volume illustration: non-photorealistic rendering of volume models. In *Proceedings of the IEEE Visualization conference*, pages 195–202, 2000.
- [4] Klaus Engel, Martin Kraus, and Thomas Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, pages 9–16. ACM Press, 2001.
- [5] Abhijeet Ghosh, Matthew Trentacoste, and Wolfgang Heidrich. Volume rendering for high dynamic range displays. In Arie E. Kaufman, Klaus Mueller, Eduard Gröller, Dieter W. Fellner, Torsten Möller, and Stephen N. Spencer, editors, *Volume Graphics*, pages 91–98. Eurographics Association, 2005.
- [6] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [7] Helwig Hauser, Lukas Mroz, Gian-Italo Bischi, and Eduard Gröller. Two-level volume rendering-fusing mip and dvr. In *Proceedings of the IEEE Visualization conference*, pages 242–252, 2000.
- [8] Gordon L. Kindlmann and James W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *IEEE Symposium On Volume Visualization*, pages 79–86, 1998.
- [9] Martin Kraus. Scale-invariant volume rendering. In *Proceedings of the IEEE Visualization conference*, pages 295–302. IEEE Computer Society, 2005.
- [10] Eric B. Lum and Kwan-Liu Ma. Lighting transfer functions using gradient aligned sampling. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 289–296, Washington, DC, USA, 2004. IEEE Computer Society.
- [11] S. Marchesin, J.-M. Dischler, and C. Mongenet. Feature enhancement using locally adaptive volume rendering. In *IEEE/EG International Symposium on Volume Graphics*, volume to appear. IEEE/EG, september 2007.
- [12] Benjamin Mora and David S. Ebert. Instant volumetric understanding with order-independent volume rendering. *Comput. Graph. Forum*, 23(3):489–498, 2004.
- [13] Benjamin Mora and David S. Ebert. Low-complexity maximum intensity projection. *ACM Trans. Graph.*, 24(4):1392–1416, 2005.
- [14] L. Mroz, E. Gröller, and A. König. Real-time maximum intensity projection. In *Data Visualization*, pages 135–144. 1999.
- [15] Takafumi Saito. Real-time previewing for volume visualization. In *VVS '94: Proceedings of the 1994 symposium on Volume visualization*, pages 99–106, 1994.
- [16] Y. Sato, N. Shiraga, and S. Nakajima. Local maximum intensity projection (lmip): A new rendering method for vascular visualization. *Journal of Computer Assisted Tomography*, Vol. 22, No. 6, November-December 1998, pages 912–917, 1998.
- [17] Petr Šereda, Anna Vilanova Bartroli, Iwo W. O. Serlie, and Frans A. Gerritsen. Visualization of boundaries in volumetric data sets using lh histograms. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):208–218, 2006.
- [18] Fan-Yin Tzeng, Eric B. Lum, and Kwan-Liu Ma. A novel interface for higher-dimensional classification of volume data. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 66, Washington, DC, USA, 2003. IEEE Computer Society.
- [19] Ivan Viola, Armin Kanitsar, and Meister Eduard Gröller. Importance-driven volume rendering. In *Proceedings of the IEEE Visualization conference*, pages 139–145, 2004.
- [20] Ivan Viola, Armin Kanitsar, and Meister Eduard Gröller. Importance-driven feature enhancement in volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):408–418, 2005.
- [21] Xiaoru Yuan, Minh X. Nguyen, Baoquan Chen, and David H. Porter. High dynamic range volume visualization. In *Proceedings of the conference on Visualization 2005*, pages 327–334. IEEE Computer Society, 2005.



**Stéphane Marchesin** is currently a post-doctoral researcher in the VIDI team at the University of California, Davis. He defended his Ph.D. in 2007 at the University of Strasbourg (France). His research interests revolve around scientific visualization, including but not restricted to volume rendering, parallel and high performance visualization.



performance visualization and volume rendering.

**Jean-Michel Dischler** is an associate professor at the University of Strasbourg (France) since 2001. He is a member of the 3D Computer Graphics Team, IGG (<https://lsit-cnrs.unistra.fr/igg-fr/>), where he is supervising the realistic rendering and scientific visualization research activities. He obtained the french research supervision diploma (HDR) from the University of Limoges, where he remained an assistant professor during five years (1996-2001). His research interests include textures, natural phenomena, real-time rendering, cultural heritage, high



**Catherine Mongenet** is professor in computer science at the University of Strasbourg (France) since 1999. She graduated from Nancy University (France) and she obtained the french research supervision diploma (HDR) from the University of Besançon (France). She is a member of the Imaging and scientific parallel computing team at the Image Sciences, Computer Sciences and Remote Sensing Laboratory. Her research interests include parallel compilers, parallel algorithms and parallel visualization.