

Obscure-based Volume Rendering Framework

M. Ruiz¹, I. Boada¹, I. Viola², S. Bruckner³, M. Feixas¹, and M. Sbert¹

¹Graphics and Imaging Laboratory, University of Girona, Spain

²Department of Informatics, University of Bergen, Norway

³Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

Abstract

Obscurances, from which ambient occlusion is a particular case, is a technology that produces natural-looking lighting effects in a faster way than global illumination. Its application in volume visualization is of special interest since it permits us to generate a high quality rendering at a low cost. In this paper, we propose an obscure-based framework that allows us to obtain realistic and illustrative volume visualizations in an interactive manner. Obscurances can include color bleeding effects without additional cost. Moreover, we obtain a saliency map from the gradient of obscurances and we show its application to enhance volume visualization and to select the most salient views.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Global illumination is a well-known technique for producing realistic scenes. However, although it might play a decisive role in 3D volume visualization since it provides visual cues that enhance data interpretation, its application is still challenging in direct volume rendering. The main limiting factor is the computational cost of simulating global effects of light in a volume, making interactive exploration difficult [Max95]. On the other hand, illustrative methods aim at creating visualizations which convey information to the viewer as opposed to physically correct light interaction. Volume illustration enhances the expressiveness of volume rendering by highlighting important features within a volume while subjugating insignificant details and rendering the result in a way that resembles an illustration [ER00]. Ideally, a volume rendering system should be able to support both realistic and illustrative renderings.

Obscurances have been introduced by Zhukov et al. [ZIK98] and Iones et al. [IKSZ03] as an efficient technique that gives perceptually similar results to global illumination with a small fraction of the computational cost. Moreover, without adding computational cost, obscurances also allow us to compute *color bleeding*, which consists in the effect that the objects around another object with intense coloration are dyed with this color [MSC03]. The obscu-

rance technique was first used in video-game environments. Its application to volume rendering, called vicinity shading, was introduced by Stewart [Ste03].

In this paper, we present an obscure-based volume rendering system that allows to obtain realistic and illustrative volume visualizations in an interactive manner. One important aspect of our work shows that obscurances are not only useful for realistic depiction but also for illustrative rendering. As obscurances can be interpreted as general information about the neighborhood of a voxel, they can be used as a bias for the generation of more expressive illustrative depictions of a data set (see Figure 1).

Saliency typically arises from contrasts between items and their neighborhood [IK01, TIR05, vdWGB06] and it is considered that the most salient voxels in a 3D data set will attract the attention of the viewer. In our approach, voxel saliency is determined by the obscure gradient, which measures the maximum variation of the obscure field. Once the saliency of the volume is obtained, we implicitly have the saliency map of any structure contained in the volume. This saliency map can be applied to viewpoint selection and to enhance visualization. This can help to discover relevant characteristics of the model otherwise unnoticed by the observer.




Figure 1: CT-human body data set rendered with the proposed obscure-based volume rendering framework. The images have been obtained by modifying interactively the transfer function and the way in which obscurances are applied to the model.

2. Background

In this section, obscurances, ambient occlusion, and related illumination models are described.

2.1. Obscurances and ambient occlusion

Zhukov et al. introduced ambient occlusion with the term obscurances [ZIK98, IKSZ03]. Roughly speaking, obscurance measures the part of the hemisphere obscured by the neighboring surfaces. For instance, a corner of a room is more obscured than the center. From the physics of light transport, obscurance expresses the lack of secondary (reflected) light rays coming to the specific parts of the scene, thus making them darker. Computation was done as a preprocess and the obscurance values were used as an ambient term during rendering. Since the obscurance computation was a property of the geometry and not of the lighting conditions, results could be combined with an arbitrary direct illumination. The method was also useful for interactive applications because the results were independent from the viewpoint. Landis detailed how ambient occlusion could be used to add realism to models [Lan02]. For a survey see [MFS08].

The obscurance O of a point p is defined as the integral

$$O(p) = \frac{1}{\pi} \int_{\Omega} \rho(d(p, \omega)) \cos \theta d\omega, \quad (1)$$

where ρ is a function of the distance $d(p, \omega)$ of the first intersection of a ray shot from point p with direction ω , p is a surface point, θ is the angle between the normal vector at p and direction ω , and the integration is over the hemisphere oriented according to the surface normal. We only consider a neighborhood of p , i.e. function ρ is set to 1 for distances greater than a maximum distance d_{max} . Therefore, the integral function $O(p)$ captures occlusion (or openness) information of the environment of point p . Considering extreme cases, an obscurance value 1 means that the point is completely open, i.e. not occluded and a value 0 means that it is completely occluded.

Ambient occlusion [Lan02] is a simplified version of the

obscurances illumination model. Ambient occlusion

$$A(p) = \frac{1}{\pi} \int_{\Omega} V(p, \omega) \cos \theta d\omega, \quad (2)$$

substitutes the ρ function in the obscurances equation (1) by the visibility function $V(p, \omega)$ that has value 0 when no geometry is visible in direction ω and 1 otherwise.

Color bleeding consists in the effect that the objects around another object with intense coloration are dyed with this color. To obtain color bleeding, Méndez et al. [MSC03] included in Equation (1) the diffuse reflectivity $R(q)$:

$$W(p) = \frac{1}{\pi} \int_{\Omega} R(q) \rho(d(p, \omega)) \cos \theta d\omega, \quad (3)$$

where q is the first point in direction ω that occludes p . When no occlusion is found within d_{max} , the average reflectivity is used. Observe that adding color bleeding to obscurances is almost free.

2.2. Volumetric shadowing

A volumetric version of the obscurances technique, called vicinity shading, was proposed by Stewart [Ste03]. Vicinity shading simulates illumination of isosurfaces by taking into account neighboring voxels. An occlusion volume is computed and stored in a shading texture that is accessed during rendering. This volume has to be re-computed each time that the rendering parameters are modified and the method does not support color bleeding. Since this first work, several models to illuminate the isosurfaces have been proposed. Wyman et al. [WPSH06] presented a method that supports the simulation of direct lighting, shadows and interreflections by storing pre-computed global illumination in an additional volume to allow viewpoint, lighting and isovalue changes. Despite the improvements achieved with these methods they still have a main limitation, they only allow to represent one of the surfaces of the volume. This limitation is overcome by Ropinski et al. [RMSD*08] and Hernell et al. [HLY07] using a local volumetric shadowing effect. Ropinski et al. compute a local histogram for each voxel from the voxel's neighbourhood, by accumulating intensities weighted by inverse squared distances. These local histograms can be combined interactively with the user defined transfer function to give an effect similar to local ambient lighting. Hernell et al. [HLY07] obtain the incident light intensity, arriving at a voxel, by integrating for each voxel and within a sphere surrounding it the attenuated transfer function density. This comes to compute, in the usual way, the visibility arriving at a voxel, using the opacities, averaged for all directions.

It is important to note at this point the twofold difference between these local volumetric shadowing effects and the classic obscurances (or ambient occlusion) used in our approach. Firstly, obscurances technique uses a ρ function (see discussion in Section 3.2) to modulate the effect of the occlusion with the distance. Secondly, obscurances compute

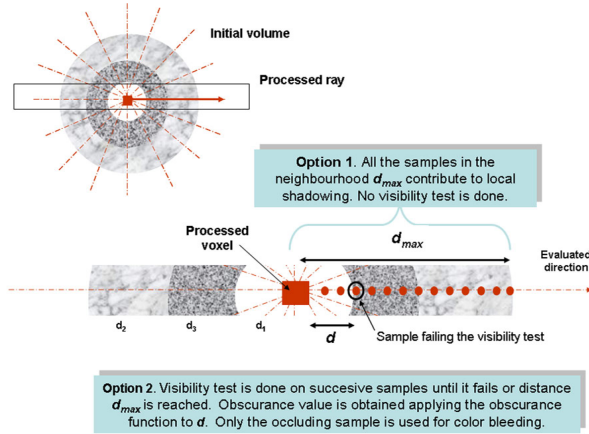


Figure 2: Local volumetric shadowing (option 1) and obscuration computation (option 2) for a volume model consisting of three concentric spheres with densities $d_1 < d_2 < d_3$.

explicitly visibility tests. This means that, although the test can be made up to a predefined maximum distance, if an occlusion is found, the rest of the environment in this direction is ignored, contrarily to local volumetric shadowing which integrates for the whole distance. In Figure 2, the essential difference between the local volumetric shadowing and the obscuration approach is illustrated. Observe that obscuration only take into account the distance from a voxel to the next occluding one, not what is in between. This is indeed different from Hernell’s algorithm, which considers the accumulated visibility of the whole environment, being nearer to the physical realism (or at least more coherent with the transfer functions).

Obscurances (and later ambient occlusion) never claimed to be physically realistic, it was introduced as a fast photo-realistic approximation to indirect illumination. Local volumetric methods have a much higher cost, with complexity proportional to the square of the number of voxels, against the complexity of obscuration computation, proportional to the number of voxels times the number of directions. Thus, on a scale of physical realism (and cost) the different strategies can be sorted into increasing order as follows: ambient occlusion (the lowest), obscuration, Hernell’s [HLY07] and Ropinski’s [RMSD*08] approaches, and global illumination (the highest).

3. Obscurances for volume rendering

In this section we go further into the obscuration-based volume rendering by testing different distance functions for computing obscuration and providing discussion on quality assumptions.

3.1. Algorithm

We take as a basic implementation of the obscuration-based volume rendering approach the one proposed by Stewart in [Ste03]. First, the volume data set is centered in a sphere built from the recursive subdivision of an icosahedron and the lines from each vertex to the center of the volume are taken as the directions to consider (12, 42, and 162 directions have been taken in our experiments). Then, for each direction, the volume is swept using Bresenham’s algorithm. This is equivalent to casting parallel rays covering all voxels. Obscuration computation for a given voxel is based on the presence (and the distance) of occluders within a certain radius along the processed direction. A visibility test compares the densities of two voxels which can be intensity values, which are independent of the transfer function, or opacity values assigned by the transfer function. In each case, we say that voxel v_i occludes v_{i-1} if the density of v_i is greater than that of v_{i-1} . To process the voxels we use a stack which stores the previously visited and yet unoccluded voxels in a density-based decreasing order. All the voxels in a ray are traversed and for each one we check if it is the nearest occluder to one of the previous unoccluded voxels (i.e. the ones stored in the stack), and in the occlusion test we check the distance. In each step, we start to check if the current voxel v_i occludes the one on the top of the stack v_s . If v_s is occluded then we can remove it from the stack so that it will not be processed anymore and continue applying the same procedure to v_{s-1} . If v_s is unoccluded, the rest of the stack voxels do not need to be processed since v_s is the voxel of the stack with lower density. Then, the next voxel of the ray is processed. This pre-computed obscuration is stored as vicinity shading values in a separate texture volume which is used during rendering.

In order to integrate color bleeding effects, we multiply the obscuration value by the color of the occluding voxel, and add an ambient color constant (in our case, white) to the unoccluded voxels. In this way, spectral obscuration are accumulated.

3.2. Analysis of ρ function

In this section, the meaning and shape of the ρ function in the obscuration definition (1) is discussed. First, this function should be a monotonically increasing function of d . Second, this function is bounded from above. This reflects the fact that normally ambient lighting of a given point is primarily affected by its neighborhood. This is especially true for scenes without bright light sources that may affect the illumination at large distances. From 0 to a determined value d_{max} , the function increases from 0 to 1, and for values greater than d_{max} the returned value is 1. This means that only a limited environment around the point p is considered and beyond this the occlusions will not be taken into account.

The shape of the ρ function is deduced from the fact that




Figure 3: Different $\rho(d)$ functions.

we are interested in what happens in the vicinity of a point (or voxel). The nearer the occlusion, the higher the influence it should have. This influence diminishes with increasing distance. This is reinforced by interpreting the obscuration model with the $\rho(d) = (1 - e^{-\tau d})$ function as the illumination at the non reflecting boundaries of a nonscattering gas with opacity τ and constant volume emittance. If we consider the occlusions of the environment as having a similar damping opacity effect over the ambient light, we should use a function as similar as possible to $\rho(d) = (1 - e^{-\tau d})$ [IKSZ03]. Some candidate functions (see Figure 3) are: (a) All-or-nothing, $\rho(d) = 0$, used in the ambient occlusion approach (it does not allow color bleeding as the contribution of a hit is 0); (b) Linear, $\rho(d) = \frac{d}{d_{max}}$, used in the vicinity shading approach; (c) Exponential, $\rho(d) = 1 - e^{-\frac{d}{d_{max}}}$; and (d) Square root, $\rho(d) = \sqrt{\frac{d}{d_{max}}}$, introduced in [MSC03]. Note that we have considered the exponential function without normalization since the normalized exponential would become very similar to the linear function.

Different data sets have been used to analyze the effect of each function in the final visualization. The obtained images are shown in Figure 4 where each column represents a different ρ : (a) all-or-nothing, (b) linear, (c) exponential, and (d) square root. In the first row, we show the behavior of the above functions (including color bleeding) on a cubic phantom model with 128^3 voxels, where except for three adjacent faces with *opacity* = 1 (left wall: green, right wall: blue, floor: white), the rest of the voxels are transparent. The effects at the corner illustrate the behaviour of the distance function. The linear function ρ (column (b)) corrects in some way the non-smoothed effect of the all-or-nothing function (column (a)), but it considers a wider than necessary environment. The exponential ρ (column (c)) produces darker values, similar to the all-or-nothing case, due to the discontinuity leap at $d = d_{max}$. The square root function (column (d)) is a good compromise, as it considers a nearer environment and the darkness decays smoothly. This is appropriate to enhance the details of the model. We also study the results of applying the different ρ functions on a CT-human body of $256 \times 256 \times 415$ voxels. The obscuration volumes and the visualizations obtained using an obscuration-based illumina-

Data	Size	12	42	162
Aneurism	$256 \times 256 \times 256$	0:36	1:56	6:56
CT-body	$256 \times 256 \times 415$	1:14	4:09	15:02

Table 1: Time cost (minutes:seconds) for computing obscurances using 12, 42, and 162 viewing directions.

tion model (see Section 5.1) are shown, respectively, in the second and third rows of Figure 4. They have been obtained considering 162 viewing directions and a maximum distance equivalent to 64 voxels. Ambient occlusion (column (a)), linear (column (b)), and exponential (column (c)) become darker, and square root (column (d)) appears less dark and more pleasant to the eye.

In the obscurances computation, d_{max} is also a key parameter to be considered since it controls the number of voxels in a determined direction that have to be taken into account to compute occlusions (see Equation 1). Therefore, if d_{max} has a high value, the probability of finding an occlusion increases, and hence the obscuration value, leading to darker images. Conversely, if d_{max} has a low value, the probability to be occluded decreases, leading to low obscuration values and hence lighter images. Figure 5 illustrates this effect in two data sets considering different d_{max} values (8, 64, and 256, respectively). As expected, we can observe how the darkness of the image increases when d_{max} increases. The other images in this paper have been computed using $d_{max} = 64$. In Figure 5, the effect of the number of directions in the obscurances values is also shown. Results for three different number of directions (12, 42, and 162) are given. Observe that although 12 directions could be considered for fast editing, we need at least 42 directions for a good quality final image. We have used the high quality obscurances given by 162 directions for the rest of the images shown in this paper. While the obscurances volume of the CT-body has been computed from its opacity (given by the transfer function), the obscurances of the aneurism have been computed from its intensity values. The time cost for computing obscurances for the CT-whole body and the aneurism is shown in Table 1. Times are given for a CPU Intel(R) Core(TM)2 Quad CPU Q6600 at 2.40GHz with 2 GB of memory. Note that, in accordance with the algorithm of Section 3.1, the time cost is proportional to the number of voxels times the number of directions. In the worst case, where densities are found in decreasing order, all the voxels in a ray would be pushed to the stack, and then each one would be popped, giving a cost proportional to the number of voxels in a ray. Thus, the cost of the algorithm is independent of d_{max} .

4. Volume saliency

The human visual system is able to reduce the amount of incoming visual data to a small but relevant amount of information for higher-level cognitive processing. Different com-




Figure 4: From top to bottom, obscurances with color bleeding for a synthetic model, obscurances for the CT-human body data set, and a rendering of this data set using an obscurance-based illumination model, all of them obtained considering different ρ functions. From left to right: (a) $\rho(d) = 0$, (b) $\rho(d) = \frac{d}{d_{max}}$, (c) $\rho(d) = 1 - e^{-\frac{d}{d_{max}}}$ and (d) $\rho(d) = \sqrt{\frac{d}{d_{max}}}$.

putational models have been proposed to interpret the selective visual attention. The biologically-inspired model of bottom-up attention of Itti et al. [IK01] permits us to understand our ability to interpret complex scenes in real time. The selection of a subset of available sensory information before further processing appears to be implemented in the form of a spatially circumscribed region of the visual field, called *focus of attention*, while some information outside the focus of attention is suppressed. This selection process is controlled by a *saliency map* which is a topographic representation of the instantaneous saliency of the visual scene and shows what humans find interesting in visual scenes.

Inspired by Itti's work, Lee et al. [LVJ05] introduced the concept of mesh saliency, a measure of regional importance for 3D meshes, computed using a center-surround mechanism that is able to identify regions that are different from

their surrounding context. Mesh saliency is captured from surface curvatures and is used in graphics applications such as mesh simplification and viewpoint selection. Feixas et al. [FSG] defined a view-based saliency of a polygon as the average information-theoretic dissimilarity between this polygon and its neighbors. In the volume rendering field, Kim et al. [KV06] presented a visual-saliency-based operator to enhance human perception of the volume data by guiding the viewer's attention to selected regions. A definition of voxel saliency is not provided and it is assumed that a saliency value is assigned to each voxel by using a user specification, eye-tracking data, or feature computation. In different works on saliency, it has been shown that attention is attracted by changes in luminance, color, curvature, texture, shape, etc. [TIR05]. That is, salient features are generally determined from the local differential structure of images and different operators such as color or luminance gradient have

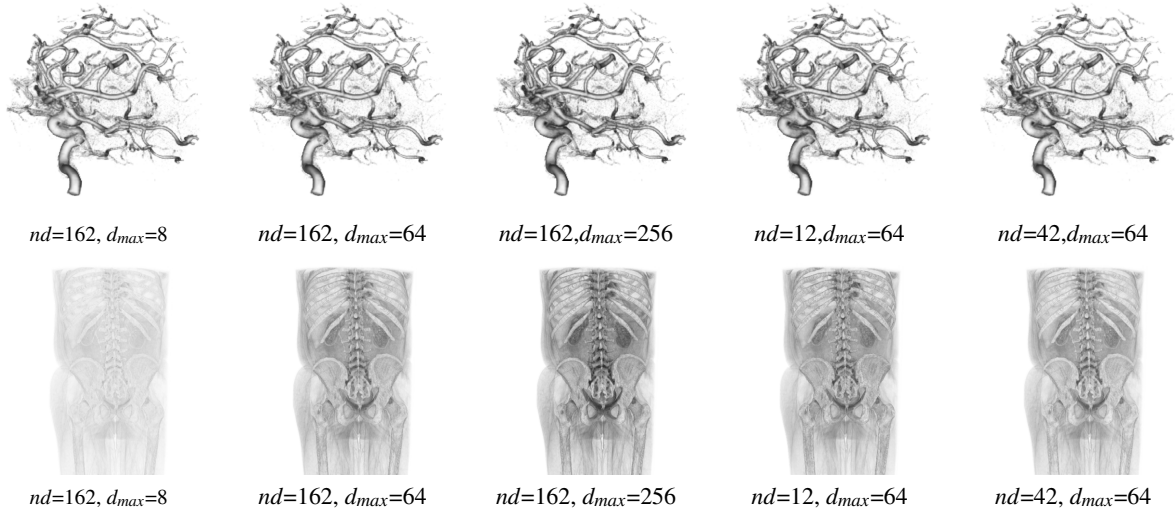


Figure 5: The obscuration volumes of the aneurism (first row) and CT-human body (second row) are visualized considering different d_{max} values and number of directions nd . The square root function has been used in all the cases. Computation times are given in Table 1.

been used [vdWGB06]. In Gonzalez et al. [GSF08], from an information theory perspective, ambient occlusion has been defined as the occlusion information associated with each polygon of the model.

In this paper, a definition of voxel saliency based on the gradient of obscuration field is proposed. Considering that obscuration represents occlusion information associated with a voxel, its variation with respect to its surround can indeed be perceptually salient, i.e. it can be considered as a salient feature of the volume. This saliency would be most noticeable at edges, occlusion variations, and corners. On the other hand, a smooth or uniform region would produce low saliency values, as is intuitively expected.

The voxel saliency is defined as the magnitude of the gradient of obscurances estimated by using the 4D linear regression method proposed in [NCKG00]:

$$S(p) = \sqrt{A^2 + B^2 + C^2}, \quad (4)$$

where voxel p is located at the origin of the coordinate system, and A , B , and C are the components of the obscuration gradient $[A, B, C]$. These components are computed as $A = \sum_k w(k)O(k)x(k)$, $B = \sum_k w(k)O(k)y(k)$, and $C = \sum_k w(k)O(k)z(k)$, where k stands for the voxels in the neighborhood centered at voxel p , $w(k)$ is the distance between voxels p and k , $O(p)$ is the obscuration of voxel p , and $x(k)$, $y(k)$, and $z(k)$ are, respectively, the x , y , and z components of the vector from voxel p to voxel k . In our experiments, the neighborhood of p is given by a cube of 5^3 voxels, since smoother results are obtained than by using a cube of 3^3 voxels as in [NCKG00]. For each data volume, the saliency has been scaled ranging from 0 to 1. Analogously to

mesh saliency [LVJ05], the gradient of obscurances is scale-dependent (i.e., the saliency value depends on the size of the neighborhood considered). We have to emphasize that our definition of saliency can be generalized to the local volumetric shadowing methods [HLY07, RMSD*08]. Figures 6(b), (d) and (f) show the color-coded saliency maps obtained for the CT-human body corresponding to Figure 6(a).

5. Applications

In this section, we describe how obscurances can be applied to volume rendering to interactively produce realistic and illustrative images. Two applications of saliency maps are also presented.

5.1. Realistic and illustrative rendering

To apply the obscurances to the visualization, we use the Blinn-Phong shading model where the color resulting from the local lighting of each voxel x is multiplied by its obscuration value:

$$I(x) = (k_d N(x) \cdot L + k_s (N(x) \cdot H)^n) O(x) \quad (5)$$

where k_d and k_s are the diffuse and specular lighting coefficients, $N(x)$ the normal of the voxel, L the light vector, H the half-angle vector between L and the direction to the viewer, n is the Phong exponent, and $O(x)$ the obscuration of voxel x which has been adjusted to the range $[0, 1]$. Figure 4 (third row) illustrates the result of applying the obscuration-based Blinn-Phong model on the CT-human body data set.

We also introduce two parameters, *low* and *high*, such that




Figure 6: (a) Original CT-body data set. (b, d, f) Color-coded (from blue to red) saliency maps corresponding to the most salient views. (c, e, g) Illustrative visualizations obtained with a saliency-based opacity modulation.

from 0 to *low* obscurances are set to 0 (making the voxel completely black), from *low* to *high* they follow a linear distribution (preserving their original value), and from *high* to 1 their value is set to 1 (thus the voxel becomes completely unobscured). Increasing the *low* threshold will turn more voxels black, and this can be used to increase the contrast of the resulting image. Decreasing the *high* threshold means that more voxels are not darkened by their obscurance, and so become brighter, thus increasing slightly the contrast too. In the limit we could set *low* and *high* to the same value to have only some voxels with obscurance 0 and others with obscurance 1. Thus, there would be voxels with their own color (modified only by local lighting if applied) and the others would be black. The user can modify *low* and *high* parameters to obtain the desired effect interactively. Figure 1 shows different renderings of the CT-body data set. Figure 1(a) shows the visualization of the model resulting from the application of obscurances with *low* = 0 and *high* = 1. Figure 1(b) has been obtained with *low* = 0.4 and *high* = 0.6, and making the skeleton transparent. Finally, in Figure 1(c) all the structures have been set to white and the obscurances assignment has been adjusted with *low* = 0.6 and *high* = 0.7.

5.2. Saliency

As a measure of importance, the volume saliency is applied to obtain the most salient views and to enhance volume visualization by modifying the transfer function according to the computed saliency.

Similar to [LVJ05], where mesh saliency was used to select the best views, a method to calculate the saliency of a viewpoint is proposed. Given the saliency of all the voxels, we can find the viewpoint which maximizes the visible saliency. The *viewpoint saliency* is defined by

$$S(v) = \sum_{p \in P} S(p)V(p), \quad (6)$$

where v is a given viewpoint, P is the set of voxels of the volume data, $S(p)$ is the saliency of voxel p , and $V(p)$ is the visibility of voxel p from v .

We also present an automated technique to enhance volume visualization by emphasizing (increasing the opacity of) the most salient voxels and de-emphasizing (reducing the opacity of) the least salient ones. So, the viewer's attention is guided towards the most salient parts of the model.

In Figure 6, (a) the original CT-body data set, (b-c) the most salient view, (d-e) the least salient view, and (f-g) the most salient view per unit area are shown. Images (c), (e) and (g) have been obtained by multiplying the opacity by the saliency. Figure 7 shows (a) the original CT-body data set and (b-c) two different renderings obtained by scaling the opacity according to the saliency values. In Figure 7(b), voxels with saliency lower than 0.2 have been made transparent and the opacity of the most salient ones has been preserved. In Figure 7(c), the voxels with saliency lower than 0.2 have been made transparent while the opacity of the most salient ones has been doubled.

6. Conclusions

In this paper, we have analyzed obscurance-based volume rendering by evaluating the main parameters involved in its computation, such as the obscurance function and the number of viewing directions. From this study, we conclude that the square root function gives better results than other analyzed functions and that 42 directions are enough to obtain obscurances of a certain quality, although for high quality results we have used 162 directions. In addition, we have introduced two new applications of obscurances. The first is a technique to obtain illustrative renderings and the second is a method to compute the saliency map as the gradient of obscurances. Saliency has been used to enhance visualization

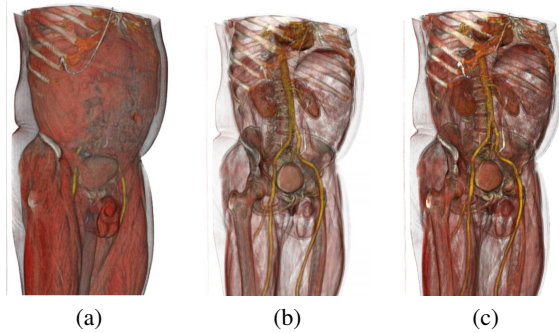


Figure 7: (a) Original CT-human body data set. (b, c) Images obtained by scaling the opacity according to the saliency values.

and to select the most salient views. All our proposals have been integrated in a common framework and tested on several volume data sets. As future work, we plan to programme a GPU version of the obscurances algorithm to obtain real-time or interactive obscurance computation.

Acknowledgements

This work has been supported by TIN2007-68066-C04-01 and TIN2007-67982-C02 of the Ministry of Education and Science (Spanish Government), by the MedViz Initiative in Bergen (medviz.uib.no), and by the Austrian Science Fund (FWF) grant no. P18322.

References

[ER00] EBERT D., RHEINGANS P.: Volume illustration: non-photorealistic rendering of volume models. In *VIS '00: Proceedings of the conference on Visualization '00* (Los Alamitos, CA, USA, 2000), IEEE Computer Society Press, pp. 195–202.

[FSG] FEIXAS M., SBERT M., GONZALEZ F.: A unified information-theoretic framework for viewpoint selection and mesh saliency. *ACM Transactions on Applied Perception*. In press.

[GSF08] GONZALEZ F., SBERT M., FEIXAS M.: Viewpoint-based ambient occlusion. *IEEE Computer Graphics and Applications* 28, 2 (2008), 44–51.

[HLY07] HERNELL F., LJUNG P., YNNERMAN A.: Efficient ambient and emissive tissue illumination using local occlusion in multiresolution volume rendering. In *Eurographics/IEEE-VGTC Symposium on Volume Graphics 2007* (2007), pp. 1–8.

[IK01] ITTI L., KOCH C.: Computational modeling of visual attention. *Nature Reviews Neuroscience* 2, 3 (2001), 194–203.

[IKSZ03] IONES A., KRUPKIN A., SBERT M., ZHUKOV S.: Fast, realistic lighting for video games. *IEEE Comput. Graph. Appl.* 23, 3 (2003), 54–64.

[KV06] KIM Y., VARSHNEY A.: Saliency-guided enhancement for volume visualization. *IEEE Trans. Vis. Comput. Graph.* 12, 5 (2006), 925–932.

[Lan02] LANDIS H.: Production-ready global illumination. In *Course 16 notes, SIGGRAPH 2002* (2002).

[LVJ05] LEE C. H., VARSHNEY A., JACOBS D. W.: Mesh saliency. *ACM Trans. Graph.* 24, 3 (2005), 659–666.

[Max95] MAX N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 1, 2 (1995), 99–108.

[MFS08] MENDEZ-FELIU A., SBERT M.: From obscurances to ambient occlusion: A survey. *The Visual Computer* (2008). In press, published online 15 February 2008.

[MSC03] MÉNDEZ A., SBERT M., CATÁ J.: Real-time obscurances with color bleeding. In *SCCG '03: Proceedings of the 19th spring conference on Computer graphics* (2003), ACM, pp. 171–176.

[NCKG00] NEUMANN L., CSÉBFALVI B., KÖNIG A., GRÖLLER E.: Gradient estimation in volume data using 4d linear regression. *Comput. Graph. Forum* 19, 3 (2000), 351–358.

[RMSD*08] ROPINSKI T., MEYER-SPRADOW J., DIEPENBROCK S., MENSMAJN J., HINRICHS K. H.: Interactive volume rendering with dynamic ambient occlusion and color bleeding. *Computer Graphics Forum (Eurographics 2008)* 27, 2 (2008), 567–576.

[Ste03] STEWART A. J.: Vicinity shading for enhanced perception of volumetric data. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)* (2003), IEEE Computer Society, pp. 355–362.

[TIR05] TSOTSOS J. K., ITTI L., REES G.: A brief and selective history of attention. In *Neurobiology of Attention*, Itti L., Rees G., Tsotsos J. K., (Eds.). Elsevier, San Diego, CA, 2005.

[vdWGB06] VAN DE WEIJER J., GEVERS T., BAGDANOV A. D.: Boosting color saliency in image feature detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 28, 1 (2006), 150–156.

[WPSH06] WYMAN C., PARKER S. G., SHIRLEY P., HANSEN C. D.: Interactive display of isosurfaces with global illumination. *IEEE Trans. Vis. Comput. Graph.* 12, 2 (2006), 186–196.

[ZIK98] ZHUKOV S., IONES A., KRONIN G.: An ambient light illumination model. In *Proceedings of Eurographics Rendering Workshop '98* (1998), Springer Wien, pp. 45–56.