

Linköping University Post Print

Local Ambient Occlusion in Direct Volume Rendering

Frida Hernell, Patric Ljung and Anders Ynnerman

N.B.: When citing this work, cite the original article.

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Frida Hernell, Patric Ljung and Anders Ynnerman, Local Ambient Occlusion in Direct Volume Rendering, 2010, IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, (16), 4, 548-559.

<http://dx.doi.org/10.1109/TVCG.2009.45>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-56687>

Local Ambient Occlusion in Direct Volume Rendering

Frida Hernell, Patric Ljung, *Member, IEEE Computer Society*, and Anders Ynnerman, *Member, IEEE*

Abstract—This paper presents a novel technique to efficiently compute illumination for Direct Volume Rendering using a local approximation of ambient occlusion to integrate the intensity of incident light for each voxel. An advantage with this local approach is that fully shadowed regions are avoided, a desirable feature in many applications of volume rendering such as medical visualization. Additional transfer function interactions are also presented, for instance, to highlight specific structures with luminous tissue effects and create an improved context for semitransparent tissues with a separate absorption control for the illumination settings. Multiresolution volume management and GPU-based computation are used to accelerate the calculations and support large data sets. The scheme yields interactive frame rates with an adaptive sampling approach for incrementally refined illumination under arbitrary transfer function changes. The illumination effects can give a better understanding of the shape and density of tissues and so has the potential to increase the diagnostic value of medical volume rendering. Since the proposed method is gradient-free, it is especially beneficial at the borders of clip planes, where gradients are undefined, and for noisy data sets.

Index Terms—Local illumination, volumetric ambient occlusion, volume rendering, medical visualization, emissive tissues, shading, shadowing.

1 INTRODUCTION

INTERACTION with patient-specific Computed Tomography (CT) data through Direct Volume Rendering (DVR) has become an invaluable tool in medical diagnosis and contributed to the spread of radiological methodologies into several areas of medical practice, such as postmortem imaging [1] and preoperative planning [2], [3], [4], [5]. To further promote the use of DVR in medicine, new techniques for enhanced perception of shapes and densities, ensuring a more confident diagnosis, are needed. One area of development that shows promise in supporting this is the inclusion of more sophisticated lighting models that provide visual cues to improve the perception of shapes and depth. Langer and Bühlhoff [6] found that ambient lighting in particular, like lighting on a cloudy day, significantly enhances the ability of humans to distinguish objects and their properties.

A commonly used lighting approximation in volume rendering applications is the Phong shading model [7], or later, the Blinn-Phong model [8]. This approach is suitable for direct volume rendering since it can be evaluated at interactive frame rates. The method does, however, make use of the normalized gradient of the volume scalar field, which limits its applicability. If, for instance, the scalar data are noisy, then the gradients become indistinct. Another problem is that the gradients are only well defined at sharp

transitions between different scalar values, which means that light values can be estimated for surfaces but not as easily for homogeneous regions. It is, therefore, more difficult to define a light approximation in semitransparent regions and at clip plane surfaces.

More advanced lighting models, like global illumination, are computationally expensive and cannot readily be used in interactive DVR. A further disadvantage of such models, in the context of medical visualization, is that regions can be shadowed by dense objects occluding the light. The skull, for example, can occlude the light such that a tumor or other feature may not be illuminated and so cannot be identified.

This paper proposes a new shading model in DVR, Local Ambient Occlusion (LAO), that considers shadowing by structures in the vicinity of each voxel. It can be used instead of, or as a complement to, diffuse shading. The method computes the incident light for each voxel by sampling a spherical neighborhood around each voxel, capturing shadows and light emissions locally. This light information is then used in the rendering pass as the ambient light term in the volume rendering integral. Since the method is gradient-free, it is also less sensitive to noise and homogeneous regions with poorly defined gradients. As an integrated part of the LAO approach, illumination from emissive tissues is also included and by using adjusted opacities in the LAO calculation, structures and features can also be made more prominent, providing more contrast, in highly transparent regions.

The main contributions of this paper are the following:

- introduction of Local Ambient Occlusion as a new shading model in direct volume rendering;
- a method for interactive transfer function-based emission that improves the visibility of user-specified data ranges;
- an adaptive sampling scheme achieving high quality light approximation while optimizing the number of samples;

• F. Hernell and A. Ynnerman are with the Department of Visual Information Technology and Applications (VITA) and the Center for Medical Image Science and Visualization (CMIV), Linköping University, Universitet Campus Norrköping 601 74, NORRKÖPING, Sweden.
E-mail: {frihe, andyn}@itn.liu.se.

• P. Ljung is with Siemens Corporate Research, 755 College Rd East, Princeton, NJ 08540-6632. E-mail: patric.ljung@siemens.com.

Manuscript received 21 Jan. 2008; revised 13 Jan. 2009; accepted 4 Mar. 2009; published online 17 Apr. 2009.

Recommended for acceptance by H.-C. Hege, D.H. Laidlaw, and R. Machiraju. For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-2008-01-0100.

Digital Object Identifier no. 10.1109/TVCG.2009.45.

- an additional transfer function for control of the amount of light absorbed by different tissue types in LAO to increase the preserved context of near-transparent tissues; and
- an alternate method enhancing shape perception at clip planes.

Despite the limited neighborhood region, estimation of LAO is computationally demanding since it must be evaluated for each voxel. A multiresolution framework is, therefore, used to reduce the amount of computation in empty or homogeneous spaces and interactivity is then maintained during the LAO calculation. Furthermore, since the ambient occlusion is view-independent, a reestimation of the integrated light intensities need only be performed when the transfer function is modified.

2 RELATED WORK

A large number of methods have been proposed for generation of realistic illumination in computer graphics [9], [10]. Computing self-shadowing in volume rendering is a computationally demanding problem and several approximation schemes have been proposed to achieve interactive frame rates. One approach is to preprocess the light computation and store the intensities in a shadow map, see Behrens and Ratering [11] and Hadwiger et al. [12], for example. This shadow map must, however, be reestimated each time the transfer function or light condition changes. If an update of the light map takes too long, it could cause problems in clinical situations, where fine-tuning of the transfer function (TF) is needed. Another approach, presented by Kniss et al. [13], estimates the shadows in parallel with the final volume rendering that is carried out using half-angle texture slicing. The half angle refers to the angle between the light direction and the view direction. This method also includes scattering and color bleeding effects. The method, however, relies on rendering based on texture slicing and supports only a single light source. A variation of this method that takes several samples in the 2D shadow buffer of each slice to create dilated shadows in the light direction was presented by Desgranges et al. [14]. The estimated shadow value of a voxel is, in this approach, the minimum value of the neighboring shadow values that also gives an effect of attenuation of light within an object. This method is gradient-free but is also still limited to a single light source.

The local volumetric shadowing effect presented in this paper is an extension of our earlier work [15] and comparable with ambient occlusion for surfaces, which is a simplification of the obscurance illumination model [16], [17], where incident light at a surface element is estimated by integrating the visibility over a hemisphere.

A number of methods have been presented for surface rendered ambient occlusion, for example, [18], [19], [20], [21]. One of these is vicinity shading, presented by Stewart [19], who modified the rendering equation presented by Kajiyama [22] to yield an illumination model that, for each surface point, estimates the light that arrives from a large number of directions. One downside of this method is long preprocessing times. Tarini et al. [18] subsequently refined this model to increase the performance. Wyman et al. [23] used a combination of Monte Carlo path tracing and a spherical harmonic representation to approximate the irradiance over a hemisphere and allow dynamic environmental lighting. A

combination of a low frequency model of the geometry and the original geometry was used by Shanmugam and Arikan [20] for a better approximation of the ambient lighting.

Methods for ambient occlusion are often referred to as “all-or-nothing” methods, which means that there is no consideration of transparency. Incident light on any voxel from a direction is blocked by any voxel having a higher value along that direction. When generalizing these methods to take participating media into account, care must be taken to reduce the complexity introduced by large regions of semitransparent voxels. A method for combining ambient occlusion computed for isosurfaces and subsurface scattering effects was recently proposed by Rezk-Salama [24].

A problem with many illumination methods is to create accurate shadowing at clipping surfaces. Wieskopf et al. [25] proposed an optical model to merge surface and volume-based illumination to achieve a consistent shading of the clipping surface. However, the accuracy of this method suffers from inaccurate reconstruction of gradients.

This paper also presents rendering with luminous effects for the purpose of highlighting tissues. Previous examples of luminous volumes can be found in rendering of astrophysical simulations as shown by Magnor et al. [26] and Kähler et al. [27]. In their work, the emittance is estimated from physical properties such as temperature and gas density. In most cases, the luminosity of the volume is precomputed and gives a static contribution to the illumination. Halo effects, as proposed by Svakhine and Ebert [28] and Bruckner and Gröller [29], are other proposed approaches for highlighting tissues in volume rendering. A further effect that the TF interactions in our presented methods bring is enhanced context in renderings of near-transparent tissues. Some methods have been proposed in previous work for context preserving visualization, for instance, by Bruckner et al. [30].

3 VOLUMETRIC AMBIENT OCCLUSION

The introduced LAO shading model generalizes the ambient lighting term typically found in shading computations. As a starting point, the traditional volume rendering integral, as expressed by Max [31], is revisited,

$$I(D) = I_0 \cdot e^{-\int_0^D \tau(t)dt} + \int_0^D g(s) \cdot e^{-\int_s^D \tau(t)dt} ds, \quad (1)$$

where the first term represents the light coming from the background I_0 , attenuated with the optical depth, the integral of the extinction coefficient τ , and the second term represents the integration of attenuated light contributions $g(s)$, for each location s , along the ray. Equation (1) represents the raycasting method used to render a volumetric image. In this paper, this process is referred as the *rendering pass*, and for clarity, the formulation of g in the rendering pass will be denoted by g_R . A closer look at the light contribution term $g_R(s)$ traditionally reveals three components,

$$g_R(s) = A(s) + k_d(\mathbf{L} \cdot \mathbf{N})c(s) + k_s(\mathbf{N} \cdot \mathbf{H})^p c(s), \quad (2)$$

where $A(s)$ is the ambient light contribution, typically represented by a constant factor, as in $k_a c(s)$. The color classification of each sample is represented by $c(s)$. The




Fig. 1. Comparison between (a) boundary light contribution and (b) volumetric light contribution. Boundary contribution causes hard shadow edges that are avoided by the volumetric light contribution.

second term represents the diffuse lighting and the third adds specular highlights based on the half-angle technique by Blinn [8].

Local Ambient Occlusion is a way to enhance the ambient term A such that shadows and light emission from local features are included. Equation (1) is used again but, instead of calculating the light reaching a pixel in the framebuffer for the final image, it is used to compute the incident light arriving at a voxel location x . LAO furthermore only considers a local neighborhood, a sphere centered at x . This serves two purposes: it both reduces the computational demand and avoids global shadows that might hide relevant information. This approach allows for interactive frame rates and yet produces sufficient shadow effects.

Equation (1) is reformulated to originate at the voxel location x , integrating the lighting contribution along rays enumerated k . The direction of integration is also changed such that the origin is at the voxel center and $D = R_\Omega$, the radius of the spherical neighborhood Ω .

$$A_k(x) = \int_a^{R_\Omega} g_A(s) \cdot e^{-\int_a^s \tau(t) dt} ds, \quad (3)$$

where a is a small offset to avoid self-occlusion of the voxel at location x and $g_A(s)$ is used to denote the special light contribution for LAO. This formulation allows incorporation of semitransparency, as compared to previous all-or-nothing approaches, and by accumulating several directions distributed over the sphere, a Local Ambient Occlusion scheme is achieved. The formulation of $g_A(s)$ is explained next.

3.1 Light Integration

As an initial approach, light could be considered to be incident on the sphere boundary only. This yields a formulation of $g_A(s)$ as

$$g_A(s) = \delta(s - R_\Omega), \quad (4)$$

where δ is the Dirac delta, ensuring only light contribution at the boundary of the neighborhood Ω . This approach yields sharp shadow boundaries as soon as any occluder intersects the sphere, see Fig. 1. To avoid this appearance, a significant number of rays must be traced to sufficiently smooth the ambient lighting. Instead, the ambient light source can be considered volumetric, adding a fractional light emission at each sample point. With the same number of integrated rays,

in this case 8, the shadows become significantly smoother, as shown in Fig. 1. The formulation of $g_A(s)$ is here

$$g_A(s) = \frac{1}{R_\Omega - a} \quad (5)$$

for a distribution of light along the ray.

Basic numerical evaluation of the integral in (3) in a front-to-back compositing scheme is simply expressed as

$$A_k(x) = \frac{1}{M} \sum_{m=1}^M g_m \prod_{i=1}^{m-1} (1 - \alpha_i), \quad (6)$$

where g_m , the discretized version of $g_A(s)$, represents the light contribution at sample point m along the ray. Since the scaling factor $1/M$ is moved outside the sum, g_m is simply equal to 1. α_i is the representation of the opacity at sample position i , according to the current TF.

The combined local ambient occlusion $A(x)$ for a voxel at x is then given by the normalized sum of all incident light rays $A_k(x)$:

$$A(x) = \frac{1}{K} \sum_{k=1}^K w_k A_k(x). \quad (7)$$

Each direction is also associated with a weight w_k to allow for directional weighting of the ambient light. Recomputing $A(x)$ is only required when the TF changes, when a different directional weighting is desired, or when the cropping of the volume is changed.

3.2 Emissive Materials and TF Variations

After establishing the basic formulation of LAO, it is possible to explore different variations of $g_A(s)$. In this work, additional TFs and variations of the original TF have been used.

An additional TF is added representing the color emission, based on user-defined tissue and material density ranges, or simply TF primitives such as trapezoids or polygon envelopes. This emission component is denoted by $c_E(s)$ and (5) is then extended to

$$g_A(s) = \frac{1}{R_\Omega - a} + c_E(s). \quad (8)$$

The emission can be implemented as a separate TF or derived from the basic TF by specifying a glow/emission factor for a TF primitive or range. This simulates single light scattering of luminous objects within the volume. For the full effect, the emission must also be included during the rendering pass. The synthetic data set used in Figs. 1 and 5 also contains subsurface layers of material, which can be made emissive with an appropriate TF selection. The result is shown in Fig. 6 and further discussed in Section 5.

In volume rendering visualizations, it can furthermore be beneficial to set a low opacity for some tissue types in order to explore inner structures while a global context is preserved. However, LAO-illuminated tissues of low densities have a homogeneous appearance with weak contrasts between shadowed and unshadowed regions since each sample point has a very low impact on the light attenuation. By allowing the user to use a separate control for the absorption in LAO, different from the one used in the rendering pass, shadow




Fig. 2. Per-fragment computations to evaluate LAO for a voxel. The packed coordinates (P) are transformed into volume coordinates (V) through an index texture. Rays are generated to sample the volume around the voxel. For each sample, the reverse index texture gives locations in the packed volume to retrieve scalar data. The final fragment intensity is the average of the incrementally processed ray intensities.

contrast can be increased for low-opacity regions. This is simply achieved by using a different lookup table for α_i in (6) or a different TF $\tau(s) = \tau_A(s)$ in (3).

4 IMPLEMENTATION

As can be noted from the equations above, LAO computation is driven by the processing of each voxel in the volume. An implementation using fully sampled regular linear volumes is, therefore, fairly straightforward. It is, however, unnecessary to compute detailed illumination in empty space or homogeneous areas, as defined by the TF. Superfluous sampling in those regions can be avoided by employing a multiresolution approach. Such an approach can remove empty blocks and select lower resolutions for homogeneous regions. In this implementation, a flat blocking multiresolution method is used, as presented by Ljung et al. [32], [33] and Ljung [34], which allows for a graceful data reduction using the TF to optimize the resulting image quality. For completeness of this paper, a summary of flat blocking data management is provided. Some additional data structures and constraints are required, described below in detail, to apply the LAO method within the multiresolution framework.

4.1 Multiresolution Processing

In a preprocessing stage, a regular linear volume is divided into small blocks of typically 16^3 voxels, each block is encoded with all levels of detail, that is, $8^3, 4^3, 2^3, 1^3$ voxels. These blocks are organized into a multiresolution representation enabling efficient loading of arbitrary LOD selection at runtime. The LOD selection is updated when the TF is changed or the volume is cropped. The resolution of each block is chosen according to an optimization process based on the blocks' content in the TF domain, i.e., the data distribution mapped through the TF, which optimizes the resulting image quality given a user- or system-defined memory limit. The scheme will prioritize blocks showing high variation in the TF domain, blocks of medium to low variation are given lower priority, and empty blocks can be skipped completely. In our implementation, the minimum and maximum LOD for nonempty blocks can be set with a separate LOD for empty blocks. Empty blocks can also be skipped.

The volume data in the LOD selection are packed into a single 3D texture, similar to the approach by Kraus and Ertl [35], but without any boundary voxel replication since an interblock interpolation technique is applied when required, see Ljung et al. [33] for details. The translation of volume coordinates V to the packed texture coordinates P

is facilitated by a forward mapping index texture. This texture has the same size as the number of blocks in three dimensions and contains the 3D offset and the size of each block. In order to do the reverse translation, a reverse mapping index texture is also required. Since a single voxel in the packed volume texture can cover an entire block, this reverse mapping index texture would have to be of the same size as the packed volume. In practice, however, it is often sufficient to restrict the minimum LOD for nonempty blocks to 4^3 voxels, and thus, allow the reverse index texture to be 64 times smaller.

4.2 LAO Pipeline

The per-fragment pipeline processing is illustrated in Fig. 2. A 3D texture is created to hold the emissive color and the ambient light of each voxel. The entire volume is then processed by mapping all the slices one by one as a rendering target of a framebuffer object. In each pass, one or more ray directions can be processed, triggered by rendering a quad over the entire framebuffer, where each pixel maps to one voxel in the mapped 3D texture. The 3D illumination texture has the same dimensions as the packed volume texture and the same multiresolution layout is used. Since the processing is driven by the multiresolution layout, the data reduction achieved by the LOD selection, thus, directly implies a performance gain in the LAO computation.

Estimation of the volumetric LAO is performed incrementally, updating one ray direction per frame and storing the accumulated result in an ambient light texture using the OpenGL blend operation. The progressive update of the ambient intensity texture, as defined by (7), is reformulated as

$$A^k = \frac{1}{k} \cdot A_k + \left(1 - \frac{1}{k}\right) \cdot A^{k-1}, \quad (9)$$

where A^k is the accumulated intensity after k ray directions, with $A^0 = 0$. The blend operation adds a fraction of the intensity A_k from ray direction k . The number of ray directions K to sample can dynamically be configured and the directions are stored in a texture together with the directional weight w_k . Ray directions are created by subdividing either a tetrahedron, icosahedron, or octahedron to different levels, depending on the desired number of rays, as in Tarini et al. [18]. These shapes generate uniform distributions of ray directions.

Once all directions are processed, the final ambient light intensity volume is used in the rendering pass. In the current implementation, the rendering is interleaved with

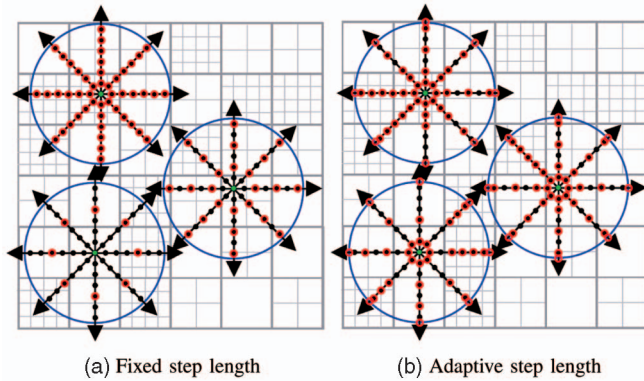


Fig. 3. Important structures can be missed when using a sampling density fixed to the LOD in the originating block (a). A more suitable sampling is performed with adaptive step lengths in (b), which optimizes the number of steps. Samples are marked in red.

the intensity texture updates, thus allowing user interaction although all ray directions may require several frames to complete.

Using the location of the fragment in the framebuffer plus the current slice z -offset, as illustrated in Fig. 2, the reverse index map $P \rightarrow V$ is used to determine the voxel location in volume coordinates. The ray integral is then computed for one or more directions. One direction per pass is commonly used to provide a higher degree of interactivity. In the integration, each sample requires a texture indirection from volume coordinates to packed texture coordinates to which end the forward mapping index texture $V \rightarrow P$ is used.

The rendering of the volume image is performed with raycasting using adaptive sampling, as presented by Ljung [34]. This technique uses the scale of each block to dynamically adapt the sampling ratio per block, and thus, effectively achieves a speedup of the order of the cube root of the data reduction ratio, the scheme is, thus, capable of skipping empty space at the front and back of the volume as well as any empty space within the volume. To provide a continuous scalar field from the multiresolution volume, an interblock interpolation technique is used [33]. This approach does not require any voxel replication from neighboring blocks and also supports arbitrary resolution differences. This interpolation scheme is applied when the user stops interacting with the volume. In many cases, the discontinuity artifacts are minor and depend on the current TF. Interblock interpolation is, however, not considered in the LAO computations, since it would require too many texture lookups. Only minor artifacts have been observed due to this limitation.

4.3 Local Ambient Occlusion Sampling

The performance of LAO computations can be improved by adjusting the sampling density along rays for voxels in lower resolution blocks. There are different ways of adjusting the sampling density and the simplest approach is to adjust the sampling rate relative to the sampling density of the underlying volume data, as illustrated in Fig. 3a. For a ray originating in a block of scale κ , the step length is adjusted with κ_{\max}/κ for all samples along that ray. The maximum resolution κ_{\max} yields the highest sampling density, given the base sampling density Δ_B . This may, however, lead to artifacts since low-resolution blocks can miss occluding features in neighboring blocks of higher




Fig. 4. A detailed illustration of sample distances along an LAO ray. The section covered by a blue area is a sample's extent Δ_s used to compute the light contribution and opacity adjustment along an LAO ray. Δ_i represent a step length computed with the adaptive sampling scheme. R_Ω is the radius of the LAO sphere and a is the initial offset to avoid self-occlusion.

resolution due to the undersampling while a neighboring block of lower resolution is highly oversampled. By employing an adaptive sampling density based on the sampled blocks resolution κ_i at sample location i along the ray, these artifacts are reduced. This ensures that features are not missed, given that a sufficient base sampling density is used. Fig. 3b illustrates this adaptive sampling strategy. A caveat is the crossing of block boundaries in which the scheme must ensure that the density is properly adjusted when entering a new block. In the stepping along the ray, the next position is set to the first discrete location, according to the base sampling density, which falls inside the next block [34]. This approach is also used in the rendering pass.

In order to maintain a consistent light contribution for the continuous volumetric light used in LAO, it is important that each ray has a unit contribution. Since an adaptive sampling scheme is used, the light contribution l_i at each sample point i must also be adaptive. This is solved by determining the light intensity based on the step sizes between the current sample locations. Fig. 4 illustrates an LAO ray with adaptive sampling and a few cases of the added light contribution for l_1 , l_3 , and l_m . The following equation describes the calculation of the light contribution for each sample point:

$$l_i = \frac{\Delta s_i}{R_\Omega - a}, \quad \Delta s_i = \begin{cases} \frac{\Delta_1}{2}, & \text{if } i = 1, \\ \frac{\Delta_{i-1} + \Delta_i}{2}, & \text{if } 1 < i < m, \\ \frac{\Delta_{i-1}}{2} + \Delta_m, & \text{if } i = m. \end{cases} \quad (10)$$

With adaptive sampling density, the opacities and colors have to be corrected as well, which is traditionally done using the opacity correction equation, below, where α' is the adjusted opacity. Δs represents the sample's extent as defined in (10). In addition, a parameter γ is introduced to further control the appearance of LAO:

$$\alpha' = 1.0 - (1.0 - \alpha)^{\frac{\Delta s}{\Delta_B}}. \quad (11)$$

The contribution of emissive light also depends on the extent of the samples and is adjusted according to

$$c'_E(s) = \frac{\Delta s}{\Delta_B} c_E(s). \quad (12)$$

4.4 Clip Plane Interaction in LAO-Illuminated Volumes

Gradient-based shading, such as diffuse illumination, is commonly used when interacting with clip planes in medical visualization. Although the diffuse shading technique is fast,




Fig. 5. A comparison of gradient-based shading and LAO. (a) shows diffuse illumination while (b) and (c) show the result of a single ray direction and 8 rays, respectively. The LAO approach generates convincing shadowing effects giving an improved 3D appearance. Frame rates are given for the rendering of the volume for a view-port of $1,024^2$. The LAO map in (b) and (c) is computed at 63 ms/ray, for a volume of $512 \times 256 \times 256$ voxels at a data reduction of 4.4:1. $R_{\Omega} = 16$ voxels.




Fig. 6. Luminous effects can be achieved by restricting the computations to be performed only to the rendering pass (a) or the LAO (b). A combination of those computations results in additional brightness of highlighted objects (c).

it has three main disadvantages. First, gradients are only well defined at sharp transitions between different scalar values, which means that gradient shading does not give realistic illumination in homogeneous regions. Second, a gradient is a local estimation that only considers the values of the closest neighbors, making it very sensitive to local variations such as noise. The third disadvantage is that gradients are undefined at transitions between the visible and invisible regions that appear when clip planes are used, unless the gradients are computed on the fly and are aware of the location of potential clip planes. Since diffuse shading is gradient-dependent, the illumination approximation is of low quality, where the gradients are undefined. By letting the clip plane affect the opacity of the samples when LAO is computed, these disadvantages are avoided and an alternative shading approach is achieved at clip planes. Outside the clip plane, the samples' opacities are reduced to zero allowing full light contribution outside the clip plane.

4.5 Extended LOD Selection for LAO Computations

The resulting resolution of a volume block, in a flat blocking multiresolution structure, depends on the block's content in the TF domain. However, since the illumination method depends on three different TFs: density, absorption, and emission, the color distortion within a block can vary depending on the choice of TF. It is, therefore, not possible to perform the LOD selection. Instead, an LOD has to be computed for each TF domain, respectively. The highest LOD of each block is chosen as the resolution for the block.

5 RESULTS

The results shown in this section examine the effect of Local Ambient Occlusion, how it performs with respect to the number of sampled ray directions, the neighborhood

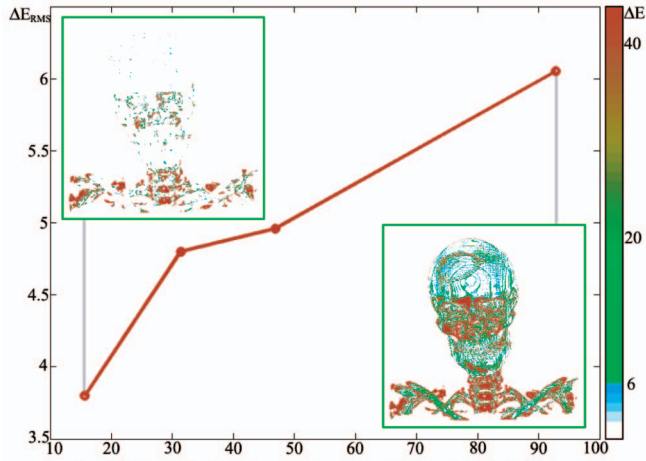
radius, and data reduction. The behavior of LAO in the presence of noise in the data set is also studied, as well as comparisons with diffuse shading and emissive tissues, and combinations of diffuse shading and LAO. All of the results have been generated using an implementation of the described methods on a standard PC equipped with a Nvidia GeForce 8800 Ultra graphics board with 768 MB of graphics memory.

The series of images in Fig. 5 show a comparison between diffuse shading and LAO for a synthetic data set with sharp gradients present between the materials and a variety of shapes from sharp corners to smooth curves. The left image shows the effect of simple diffuse lighting, providing cues about object shape. A single ray direction is used in the middle image, showing the local shadowing effect. The right image shows the effect of eight ray directions, already showing a smooth shadow approximation of local ambient occlusion.

The inclusion of emissive TF components is shown in Fig. 6. Fig. 6a considers only the effect of emissive material in the rendering pass. The emitted light does not, therefore, illuminate its environment. In Fig. 6b, emission is included only in LAO, so shadows are reduced in regions, where emissive materials are nearby. Fig. 6c shows the combined effect of emission in both LAO and the rendering pass.

5.1 Illumination Quality

A series of comparisons have been conducted to study the effect of LAO under different conditions. In order to quantify the differences, an error measurement has been used, based on a perceptually adapted color error ΔE in the CIE 1976 $L^*u^*v^*$ color space (CIELUV) [32], [36]. The pixelwise difference ΔE is designed to provide a Just Noticeable Difference (JND) at 1.0. The root mean square error ΔE_{RMS} only gives an overall error, and thus, an



(a) Graph of data reduction (x-axis) versus ΔE_{RMS} (y-axis). The color-map used in the color coded images is shown to the right.



(b) 8.5:1 data reduction (c) 15.7:1 data reduction (d) 92.8:1 data reduction

Fig. 7. A graph showing the ΔE_{RMS} in images computed and rendered with varying data reductions is provided in (a). The rendered images are compared to an image based on a data reduction of 8.5:1 (b). Two data reductions, 15.7:1 and 92.8:1, are chosen from the graph and the corresponding color-mapped difference images are shown in the green squares. Their original images are shown in (c) and (d), respectively.

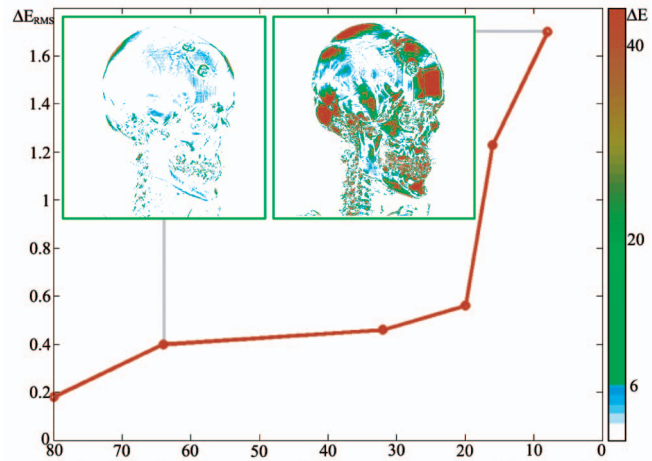
additional measure ΔE_6 shows the ratio of pixels with difference greater than 6.0.

Different levels of data reduction can be achieved based on the memory constraints assigned to the multiresolution data management. The ΔE_{RMS} , from a number of images rendered with different data reductions, compared to an image based on a reduction of 8.5:1, Fig. 7b, are plotted in the graph, Fig. 7a, to show the impact of the illumination approximation. The error measure cannot solely capture the errors from LAO but also includes errors in the rendering due to raycasting the data set at the same data reduction. However, considering that 92.8:1 is a significant data reduction, with $\Delta E_{RMS} = 6.0$ and $\Delta E_6 = 18.0$ percent, the errors are quite small. In comparison, the image generated at a data reduction of 15.7:1 has $\Delta E_{RMS} = 3.8$ and $\Delta E_6 = 4.7$ percent.

The difference when varying the number of ray directions in the LAO computation has also been measured. An image rendered using 128 directions is considered to be the “gold standard” (Fig. 8b) and compared with images based on 80, 64, 32, 20, 16, and 8 directions. With 16 or 8 ray directions, the differences sharply increase, although the appearance of LAO is still maintained.

5.2 Performance of Local Ambient Occlusion

The illumination model has also been tested for performance improvement as a function of data reduction. Fig. 9 shows the performance in frames per second (FPS) at various levels of data reduction for a volume with a native resolution of 512^3 voxels. In this case, time is measured for an LAO map



(a) Graph of number of rays (x-axis) versus ΔE_{RMS} (y-axis). The color-map used in the color coded images is shown to the right.



(b) 128 rays (c) 64 rays (d) 8 rays

Fig. 8. The ambient effect of LAO decreases when too few directions are used. The graph in (a) shows the ΔE_{RMS} of images rendered with varying numbers of directions compared to an image rendered with 128 directions. Difference images of 64 and 8 rays are provided in the green squares of the graph and their original images are provided in (c) and (d), respectively.

computed in one direction, which corresponds to the calculations needed during one incremental LAO update, independent of the number of rays in which Ω is discretized. Two graphs are shown in the performance chart, which refer to the fixed sampling rate (red) and the adaptive sampling




Fig. 9. Graph of the LAO computation time per ray versus data reduction for the data set (512^3 voxels) and TF settings as in Fig. 10. Performance is measured for fixed sampling (dotted red) and adaptive sampling (solid blue). The performance of the LAO computation is slightly superlinear with respect to the data reduction. The performance of the adaptive sampling is slightly reduced due to the higher quality and decreased parallelism. $R_\Omega = 16$ voxels and $M = 15$ are used for the highest resolution in the fixed sampling rate.




Fig. 10. The density of the skull is more apparent with LAO (b) compared to diffuse shading (a). However, high-frequency surface details are lost, and in some situations, it can be beneficial to combine the methods in order to explore fine structures, see Fig. 16.

rate (blue). The performance increase of the LAO map creation is slightly superlinear in the level of data reduction since the reduction increases the number of low-resolution blocks present in the volume with two contributing effects: the first being that the number of voxels is decreased and so less occlusion calculations are required and the second that low-resolution blocks permit a longer sample step and so the average time per occlusion calculation is reduced. The performance change due to using adaptive and fixed step length is limited, but it depends on the mixture of LOD blocks. One negative effect of adaptive sampling is an increased execution time due to the need to process different numbers of samples for neighboring voxels. However, the resulting illumination has fewer artifacts when LAO computations are performed using adaptive sampling.

5.3 Benefits and Limitations

This section discusses the benefits and limitations of the LAO method in the context of the following problem areas:

1. understanding of shapes and densities;
2. local relative position of complex structures;




Fig. 11. Close-ups of the T in images (a-d), and (h-k) of Fig. 14. The shading and perception of the shapes in the diffusely illuminated T (upper row) are much more sensitive to noise than the LAO-illuminated T (lower row). The rectangular shape of the T stem is still perceived in all images for LAO, but almost lost at even the lowest noise level for diffuse illumination.

3. artifacts in noisy data sets;
4. illumination at cross sections; and
5. tissue separation using emission.

The two images in Fig. 12 show how the LAO emphasizes the complex shape of the blood vessels and their position relative to the skull surface. LAO also enhances the visibility of the skull structure itself. A limitation, and at the same time a flexibility, of the method is the user-specified size of the LAO region, which determines at what scale shadowing effects will be included. With an increase of the LAO sphere, more neighboring structures are included and the contrast in shadows is enhanced, as can be seen in Fig. 13. R_{Ω} is, here, the radius of the surrounding LAO sphere. Another benefit of the method is the possibility to emphasize density of material using light propagation, such as the skull in (Fig. 10).

LAO is less sensitive to noisy data sets compared with gradient-based methods, such as diffuse illumination. Gaussian white noise is added to the synthetic data set, with a mean of 0 and variances of 0.0001, 0.001, and 0.01. Close-up views are shown in Fig. 11 and the full views are shown in Fig. 14. The shape of the "T" is almost lost in the diffuse illuminated




Fig. 12. Example images showing the enhanced 3D structure made clear through the LAO method. $R_{\Omega} = 32$ voxels.




Fig. 13. A large radius of the surrounding sphere Ω includes more structures in the vicinity of each voxel in LAO computations. The shadows, therefore, become darker and more defined. Measurements of computing one ray in LAO are shown within parenthesis.

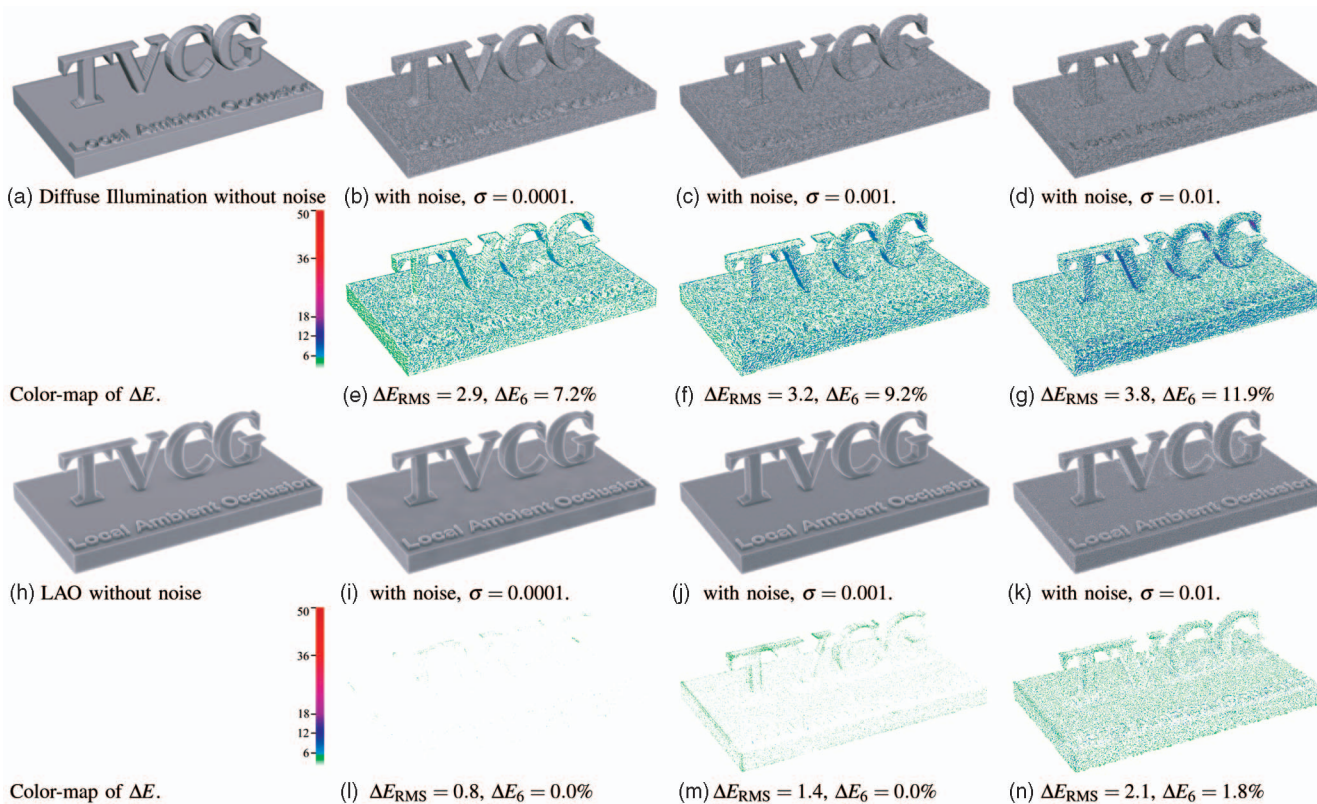


Fig. 14. Comparing diffuse shading and LAO with increasing noise in the data. (a) and (h) are renderings of a synthetic data set illuminated with diffuse shading and LAO, respectively. Gaussian noise has been added to the data set with a mean of zero and different variance σ . (b)-(d) are illuminated with diffuse shading and (i)-(k) are illuminated with LAO. Errors are provided in (e)-(g) and (l)-(n) for the noisy data sets relative to the image of their respective non-noisy data set using the same shading technique. Errors in LAO primarily appear on edges in data sets with low variance.

images while it is preserved in the LAO-illuminated images. The images in Figs. 14b, 14c, and 14d are rendered with diffuse shading and compared to Fig. 14a, while images in Figs. 14i, 14j, and 14k are illuminated with LAO and compared to Fig. 14h. Errors are computed with the method described above and the error images are found in Figs. 14e, 14f, 14g and Fig. 14l, 14m, 14n, respectively. In the LAO-rendered images, it is primarily the edges that are affected when noise of low variance is added.

When the volume has a high signal-to-noise-ratio (SNR), it can also be valuable to combine LAO with diffuse shading to explore fine structures, as can be seen in Fig. 16c. Intensifying structures with gradient-dependent shading in volumes with

low SNR, however, becomes more disturbing than helpful. It is important to note that the LAO method implicitly filters the light information with a kernel that has the extent of the LAO sphere. In cases when a trained eye is looking for details in the data, diffuse shading (despite noise) may be more relevant than the improved overall shading quality. As an example of shading of noisy data, an MRI data set is shown in Fig. 15. Interaction with clip planes is an important feature in medical DVR and LAO has the potential to improve the quality of the rendering of the cross section. The boundaries of the active volume often have poorly defined gradients and gradient-based illumination is, therefore, inappropriate. Using LAO in the illumination of the MRI data set, as shown in Fig. 15,

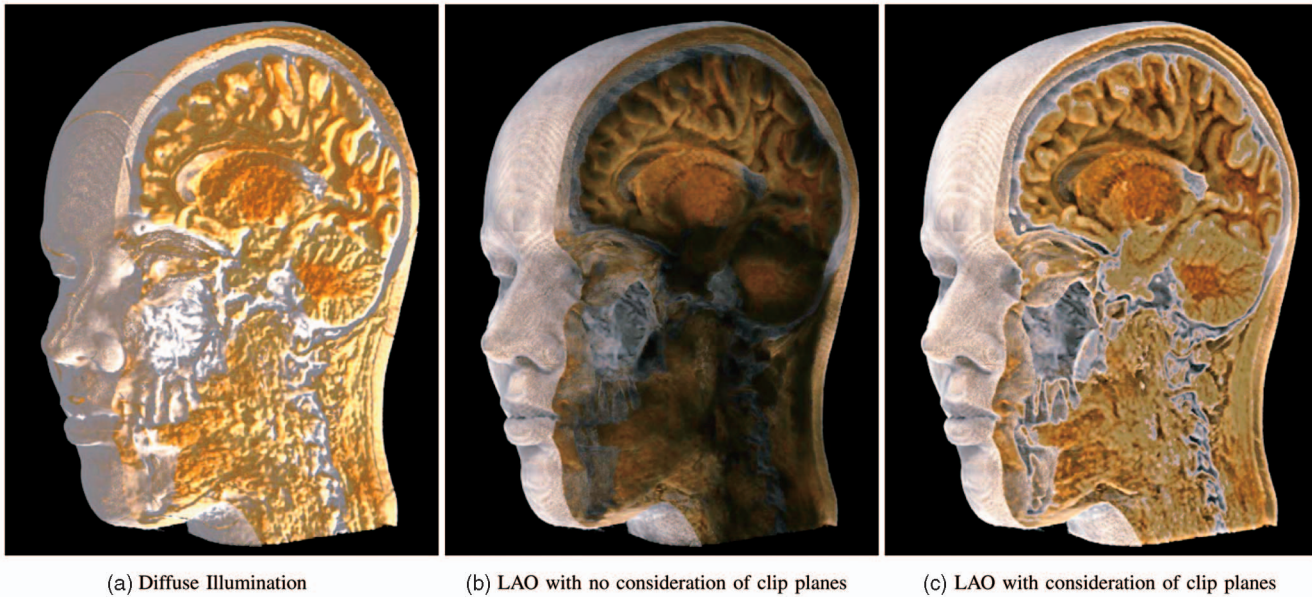


Fig. 15. A magnetic resonance (MR) volume with different illumination methods. The undefined gradients at the clipped boundary results in misleading illumination when using diffuse shading (a). Shapes are more easily perceived in the LAO-illuminated volume (c). (b) illustrates the illumination without consideration of clip planes whereas (c) is computed with additional light at the clipped boundary.




Fig. 16. Example images showing the enhanced information from the emissive materials. The bullet and fragments are clearly visible in the abdomen. The effect of the LAO in revealing the bone structure is also very clear. A combination of LAO and diffuse shading intensifies fine details, for example, the structure of the pelvis.

significantly enhances the visibility of shapes. Radiologists have verified the potential value of the LAO method for clip plane rendering.

The effects of the emissive material within a medical volume can be seen in Fig. 16, where the rendering of a virtual autopsy case is shown. In these images, the bullet and the many fragments, which are crucial to the forensic pathologist, are not as evident with diffuse shading as they are with LAO and emissive tissues. The middle image clearly shows these foreign bodies as well as revealing the bone structure more clearly. The emissive material approach compares well with the dual TF technique introduced in [1] to highlight similar foreign bodies. Combining LAO and diffuse shading (right) can, in some applications, intensify the perception of shapes, which indicates that diffuse shading has benefits over LAO in cases when there is a well-defined transition with a sharp gradient in the data.

It can be difficult to perceive a global context in applications, where semitransparent tissues are used to complement other important structures in the volume. Using a separate absorption TF in the LAO computations opens possibilities for intensifying shadowing for semitrans-

parent tissues by increasing the absorption. An example is shown in Fig. 17, where the right image uses a separate absorption setting for the soft tissue of a lion. The shape of the skin appears much clearer in the right image when compared with the left. Improved context is only noticeable when the absorption of the tissue used in the rendering pass is lower than the LAO absorption.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an efficient method for inclusion of local ambient and emissive lighting for volume rendering applications. Restricting the ambient occlusion to a local neighborhood, combined with a multiresolution framework and adaptive sampling, enables rendering at interactive speeds, a delay only being introduced by the recalculation of the occlusion volume when transfer functions are changed, and then providing a progressive refinement that allows uninterrupted user interaction.

The results show that the LAO method has properties that can improve DVR when used on its own or when combined




Fig. 17. A CT scan of a lion ($512 \times 512 \times 1,100$ voxels) rendered using the same TF for LAO and rendering (left and top TF). Using a specific TF, with increased opacity for LAO (right and bottom TF), creates a more dynamic shading effect. The dotted lines refer to the absorption TF in the LAO computations while the height of the opaque trapezoids refers to the absorption in the rendering.

with other shading models. In particular, the salience of local features in the volume can be improved, especially when the depth order of structures is crucial. Furthermore, the method is less sensitive to noisy data than gradient-based shading approaches, which can be exploited to improve the visual impression of data from, for instance, MRI scans.

The developed methods have been tested by radiologists who considered LAO to be promising for diagnosis and surgical planning. Further evaluation through a comprehensive user study is, however, needed to establish the improved diagnostic value of LAO in medicine, and for which general cases, it improves the quality of volume-rendered images. We foresee many possible further developments. Directional ambient occlusion could yield very distinct shadows, while allowing the user to interactively alter the lighting could further improve the depth perception. This would require faster calculation of the illumination parameters after changes in the light direction. Another area of interest is the relation of emission with tissue classification methods to further improve detection of regions of interest such as blood vessels or tumors.

ACKNOWLEDGMENTS

This work has been funded by the Swedish Research Council grant 621-2004-3829 and the Strategic Research Center MOVIII, founded by the Swedish Foundation for Strategic Research, SSF. The medical data sets used are provided by Siemens and the Center for Medical Image Science and Visualization (CMIV). The authors would like to thank the coworkers at the Division for Visual Information Technology and Applications.

REFERENCES

- [1] P. Ljung, C. Winskog, A. Persson, C. Lundström, and A. Ynnerman, "Full Body Virtual Autopsies Using a State-of-the-Art Volume Rendering Pipeline," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 869-876, Sept./Oct. 2006.
- [2] L. Serra, R.A. Kockrol, C.G. Guan, N. Hern, E.C.K. Lee, Y.H. Lee, C. Chan, and W.L. Nowinski, "Multimodal Volume-Based Tumor Neurosurgery Planning in the Virtual Workbench," *Proc. Int'l Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI '98)*, vol. 1496, pp. 1007-1015, 1998.
- [3] P. Jannin, O.J. Fleig, E. Seigneuret, X. Mor, M. Raimbault, and R.C. France, "Multimodal and Multi Informational Neuro-Navigation," *Proc. Conf. Computer Assisted Radiology and Surgery (CARS)*, pp. 167-172, 2000.
- [4] A. Neubauer, L. Mroz, S. Wolfsberger, R. Wegenkittl, M.-T. Forster, and K. Buhler, "Steps—An Application for Simulation of Transsphenoidal Endonasal Pituitary Surgery," *Proc. IEEE Conf. Visualization 2004*, pp. 513-520, Oct. 2004.
- [5] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Bühler, "High-Quality Multimodal Volume Rendering for Preoperative Planning of Neurosurgical Interventions," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1696-1703, Nov./Dec. 2007.
- [6] M.S. Langer and H.H. Bühlhoff, "Perception of Shape from Shading on a Cloudy Day," Technical Report 73, Max-Planck Institut für biologische Kybernetik, Oct. 1999.
- [7] B.T. Phong, "Illumination for Computer-Generated Images," PhD dissertation, The Univ. of Utah, 1973.
- [8] J.F. Blinn, "Models of Light Reflection for Computer Synthesized Pictures," *ACM SIGGRAPH Computer Graphics*, vol. 11, no. 2, pp. 192-198, 1977.
- [9] H.W. Jensen, *Realistic Image Synthesis Using Photon Mapping*, A.K. Peters, Ltd., 2001.
- [10] E. Cerezo, F. Perez-Cazorla, X. Pueyo, F. Seron, and F. Sillion, "A Survey on Participating Media Rendering Techniques," *The Visual Computer*, <http://artis.inrialpes.fr/Publications/2005/CPSS05>, 2005.
- [11] U. Behrens and R. Ratering, "Adding Shadows to a Texture-Based Volume Renderer," *Proc. IEEE Symp. Volume Visualization*, pp. 39-46, 1998.
- [12] M. Hadwiger, A. Kratz, C. Sigg, and K. Bühler, "GPU-Accelerated Deep Shadow Maps for Direct Volume Rendering," *Proc. ACM Eurographics/SIGGRAPH*, pp. 49-52, 2006.
- [13] J. Kniss, S. Premoze, C. Hansen, and D. Ebert, "Interactive Translucent Volume Rendering and Procedural Modeling," *Proc. IEEE Conf. Visualization*, pp. 109-116, 2002.
- [14] P. Desgranges, K. Engel, and G. Paladini, "Gradient-Free Shading: A New Method for Realistic Interactive Volume Rendering," *Proc. Conf. Vision, Modelling, and Visualization*, Nov. 2005.
- [15] F. Hernell, P. Ljung, and A. Ynnerman, "Efficient Ambient and Emissive Tissue Illumination Using Local Occlusion in Multi-resolution Volume Rendering," *Proc. Eurographics/IEEE-VGTC Symp. Volume Graphics*, 2007.
- [16] S. Zhukov, A. Inoes, and G. Kronin, "An Ambient Light Illumination Model," *Rendering Techniques*, G. Drettakis and N. Max, eds., pp. 45-56, Springer-Verlag Wien, 1998.
- [17] A. Iones, A. Krupkin, M. Sbert, and S. Zhukov, "Fast, Realistic Lighting for Video Games," *IEEE Computer Graphics and Applications*, vol. 23, no. 3, pp. 54-64, May 2003.
- [18] M. Tarini, P. Cignoni, and C. Montani, "Ambient Occlusion and Edge Cueing to Enhance Real Time Molecular Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1237-1244, Sept./Oct. 2006.
- [19] A.J. Stewart, "Vicinity Shading for Enhanced Perception of Volumetric Data," *Proc. IEEE Conf. Visualization*, pp. 355-362, 2003.

- [20] P. Shanmugam and O. Arikian, "Hardware Accelerated Ambient Occlusion Techniques on GPUs," *Proc. Conf. Interactive 3D Graphics and Games*, pp. 73-80, 2007.
- [21] M. Sattler, R. Sarlette, G. Zachmann, and R. Klein, "Hardware-Accelerated Ambient Occlusion Computation," *Proc. Conf. Vision, Modeling, and Visualization*, pp. 119-135, <http://www.gabrielzachmann.org/>, Nov. 2004.
- [22] J.T. Kajiya, "The Rendering Equation," *Proc. ACM SIGGRAPH*, vol. 20, no. 4, pp. 143-150, 1986.
- [23] C. Wyman, S. Parker, P. Shirley, and C. Hansen, "Interactive Display of Isosurfaces with Global Illumination," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 2, pp. 186-196, Mar./Apr. 2006.
- [24] C. Rezk-Salama, "GPU-Based Monte-Carlo Volume Raycasting," *Proc. Conf. Pacific Graphics*, 2007.
- [25] D. Weiskopf, K. Engel, and T. Ertl, "Interactive Clipping Techniques for Texture-Based Volume Visualization and Volume Shading," *IEEE Trans. Visualization and Computer Graphics*, vol. 9, no. 3, pp. 298-312, July-Sept. 2003.
- [26] M.A. Magnor, K. Hildebrand, A. Lintu, and A.J. Hanson, "Reflection Nebula Visualization," *Proc. IEEE Conf. Visualization*, pp. 255-262, 2005.
- [27] R. Kähler, J. Wise, T. Abel, and H.-C. Hege, "GPU-Assisted Raycasting for Cosmological Adaptive Mesh Refinement Simulations," *Proc. Conf. Volume Graphics*, pp. 103-110, 2006.
- [28] N.A. Svakhine and D.S. Ebert, "Interactive Volume Illustration and Feature Halos," *Proc. 11th Pacific Conf. Computer Graphics and Applications (PG '03)*, p. 347, 2003.
- [29] S. Bruckner and E. Gröller, "Enhancing Depth-Perception with Flexible Volumetric Halos," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1344-1351, Nov./Dec. 2007.
- [30] S. Bruckner, S. Grimm, A. Kanitsar, and M.E. Gröller, "Illustrative Context-Preserving Volume Rendering," *Proc. EuroVis Conf.*, pp. 69-76, May 2005.
- [31] N. Max, "Optical Models for Direct Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 2, pp. 99-108, June 1995.
- [32] P. Ljung, C. Lundström, A. Ynnerman, and K. Museth, "Transfer Function Based Adaptive Decompression for Volume Rendering of Large Medical Data Sets," *Proc. IEEE Conf. Volume Visualization and Graphics*, pp. 25-32, 2004.
- [33] P. Ljung, C. Lundström, and A. Ynnerman, "Multiresolution Interblock Interpolation in Direct Volume Rendering," *Proc. IEEE Eurographics/Visualization Conf.*, pp. 259-266, 2006.
- [34] P. Ljung, "Adaptive Sampling in Single Pass, GPU-Based Raycasting of Multiresolution Volumes," *Proc. Eurographics/IEEE Int'l Workshop Volume Graphics*, pp. 39-46, 2006.
- [35] M. Kraus and T. Ertl, "Adaptive Texture Maps," *Proc. Eurographics/ACM SIGGRAPH*, pp. 7-15, 2002.
- [36] M.D. Fairchild, *Color Appearance Models*, Addison Wesley Longman, Inc., 1998.



Frida Hernell received the MS degree in media technology in 2004 from Linköping University, Sweden. Since 2005, she has been working toward the PhD degree in scientific visualization at the Visual Information Technology and Applications (VITA) and CMIV. From 2004 to 2005, she worked as a research engineer at the Center for Medical Image Science and Visualization (CMIV) at the University Hospital in Linköping. Her research interest includes volume rendering, illumination, and interactive visualization.



Patric Ljung received the MS degree in information technology and the PhD degree in scientific visualization from Linköping University, Sweden, in 2000 and 2006, respectively. Between 2000 and 2007, he was a faculty staff member at the Norrköping Visualization and Interaction Studio (NVIS) and the Visual Information Technology and Applications (VITA) Group, also collaborating with CMIV at Linköping University. Since 2007, he has been a research scientist in the Imaging and Visualization Department at Siemens Corporate Research in Princeton, New Jersey. From 1989 to 1995, he worked as a software engineer with embedded and telecom systems. His research interest includes interactive visualization of large data sets, volume rendering, graphics system design, and software design and engineering. He is a member of the IEEE Computer Society, Eurographics, and the ACM/SIGGRAPH.



Anders Ynnerman received the BSc and PhD degrees in physics respectively from the University of Lund in 1986 and Gothenburg University in 1992. He has directed the national supercomputing organizations NSC and SNIC. Since 1999, he has been a professor of scientific visualization at Linköping University, and in 2000, he founded the Norrköping Visualization and Interaction Studio (NVIS). His current primary research interest lies in the area of visualization of large data sets and multimodal interaction techniques. He is a member of the IEEE, the IEEE Computer Society, Eurographics, and the ACM.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.