

# Volumetric Ambient Occlusion for Volumetric Models

Marc Ruiz · László Szirmay-Kalos · Tamás Umenhoffer · Imma Boada · Miquel Feixas · Mateu Sbert

the date of receipt and acceptance should be inserted later

**Abstract** This paper presents new algorithms to compute ambient occlusion for volumetric data. Ambient occlusion is used in video-games and film animation to mimic the indirect lighting of global illumination. We extend a novel interpretation of ambient occlusion to volumetric models, that measures how big portion of the tangent sphere of a surface belongs to the set of occluded points, and propose statistically robust estimates for the ambient occlusion value. The data needed by this estimate can be obtained by separable filtering of the voxel array. As ambient occlusion is meant to obtain global illumination effects, it can provide decisive clues in interpreting the data. The new algorithms obtain smooth shading and can be computed at interactive rates, being thus appropriate for dynamic models exploration.

**Keywords** Ambient Occlusion · Obscurances · Volume Rendering

---

M. Ruiz · I. Boada · M. Feixas · M. Sbert  
Graphics and Imaging Laboratory, Universitat de Girona, Campus  
Montilivi, E-17071, Girona, Spain  
E-mail: mruiz@ima.udg.cat

I. Boada  
E-mail: imma.boada@udg.edu

M. Feixas  
E-mail: miquel.feixas@udg.edu

M. Sbert  
E-mail: mateu.sbert@udg.edu

L. Szirmay-Kalos · T. Umenhoffer  
Budapest University of Technology and Economics  
E-mail: szirmay@iit.bme.hu

T. Umenhoffer  
E-mail: umitomi@freemail.hu

## 1 Introduction

Global illumination is a well-known technique for producing realistic scenes. It might also play a decisive role in 3D volume visualization since it provides visual clues that enhance data interpretation. However 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 difficult interactive volume explorations.

Obscurances or ambient occlusion (AO) was introduced by Zhukov et al. [17], Iones et al. [4], and Landis [6] as an efficient technique that gives perceptually similar results to global illumination with a small fraction of the computational cost. The obscurance/ambient occlusion technique was first used in video-game environments. Its application to volume rendering, called vicinity shading, was introduced by Stewart [13].

This paper presents the volumetric ambient occlusion for volumetric datasets. The proposed approach is based on the volumetric ambient occlusion model [14, 10], where ambient occlusion is interpreted as how big portion of the tangent sphere of a surface belongs to the set of occluded points. The integrand of the new formula has low variation, thus can be estimated accurately with a few samples. The algorithm can effectively be used in real-time systems and games to cheaply approximate global illumination effects. In our approach we extend this interpretation to compute the ambient occlusion of a voxel. The proposed approach runs in interactive time on current GPUs due to the use of separable filtering. It does not require any pre-processing and thus can be applied to general dynamic models, providing high quality smooth shading.

## 2 Previous work

In this section we describe the obscurance and related illumination models. We also review previous work on volumetric ambient occlusion.

### 2.1 Obscurances and ambient occlusion

Zhukov et al. introduced ambient occlusion with the term obscurances [17, 4]. Roughly speaking, obscurance measures the part of the hemisphere obscured by the neighbor surfaces. 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 view independent. Landis detailed how ambient occlusion could be used to add realism to models [6] and Méndez et al. extended the use of obscurances to ray tracing [7]. For a survey see [8].

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

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

where  $\rho$  is a function of the distance  $d(\mathbf{x}, \omega)$  of the first intersection of a ray shot from point  $\mathbf{x}$  with direction  $\omega$ ,  $\mathbf{x}$  is a surface point,  $\theta$  is the angle between the normal vector at  $\mathbf{x}$  and direction  $\omega$ , and the integration is over the hemisphere oriented according to the surface normal. We only consider a neighborhood of  $\mathbf{x}$ , i.e. function  $\rho$  is set to 1 for distances greater than a maximum distance  $R$ . Therefore, the integral function  $O(\mathbf{x})$  captures occlusion (or openness) information of the environment of point  $\mathbf{x}$ . Considering the 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 [6] is a simplified version of the obscurances illumination model. Ambient occlusion

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

substitutes the  $\rho$  function in the obscurances Equation 1 by the visibility function  $V(\mathbf{x}, \omega)$  that has value 0 when no geometry is visible in direction  $\omega$  and 1 otherwise.

The evaluation of the directional integral in the ambient occlusion formula (Equation 1) requires rays to be traced in many directions, which is rather costly and needs complex GPU shaders. Under the hypothesis of at most a single intersection of a traced ray within the radius  $R$  and taking

$\rho = d^3$ , the directional integral of Equation 1 can be transformed to a volumetric one (Equation 3) and the expensive ray tracing operation is substituted by a simple containment test [14]. The domain of integration, the hemisphere above the surface, becomes now the tangent sphere.

Thus, the volumetric ambient occlusion is the relative volume of the unoccluded part of the tangent sphere  $S$ . Formally, the volumetric ambient occlusion function is defined as:

$$O(\mathbf{x}) = \frac{\int_S \mathcal{I}(\mathbf{p}) d\mathbf{p}}{|S|}, \quad (3)$$

where  $\mathcal{I}(\mathbf{p})$  is 1 if point  $\mathbf{p}$  is in the unoccluded region and zero otherwise, and  $|S| = 4(R/2)^3\pi/3$  is the volume of the tangent sphere, which makes sure that the volumetric ambient occlusion is also in  $[0, 1]$ .

### 2.2 Volumetric shadowing

A volumetric version of the obscurances technique, called vicinity shading, was proposed by Stewart [13]. 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 does not support color bleeding. Since this first work, several models to illuminate the isosurfaces have been proposed.

Desgranges and Engel [1] proposed a method which combines ambient occlusion volumes terms from a plurality of different filtered volumes. Recomputation is also required whenever the transfer function is changed.

Wyman et al. [16] 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. [11] and Hernell et al. [3], using a local volumetric shadowing effect. Ropinski et al. compute for each voxel a local histogram 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. [3] 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.

Ruiz et al. [12] presented an obscurance-based volume rendering system that allows to obtain realistic and illustrative volume visualizations in an interactive manner. They interpreted obscurances as general information about the neighborhood of a voxel, and used them as a bias for the generation of more expressive illustrative depictions of a dataset.


Jainek et al. [5] combined ambient occlusion calculations with illustrative display styles to enhance the clarity of the visual output in brain studies.

Penner and Mitchell [10] computed an approximation of ambient occlusion using the same formula as Equation 3, where the domain of integration was the neighborhood around a surface point (they take both hemisphere and tangent sphere), based on the heuristic assumption that the percentage of occluding voxels surrounding the surface provides a good approximation to the percentage of rays that would be occluded while traversing through the same space. They used one-tailed Chebychev’s inequality to compute them as *variance occlusion maps*, but stated in the conclusions that, although Chebychev’s inequality was a quick method to approximate complex lighting effects, better CDF approximations should be used.

This paper is related to this approach and proposes more robust estimates for ambient occlusion.

### 3 New approaches for volumetric ambient occlusion

We propose two methods to approximate the ambient occlusion from the Cumulative Distribution Function (CDF) of the density values. The parameters of the estimation are computed by separable filters. Following the basic assumptions in [14], we will keep to the tangent sphere for a more robust interpretation of the results as Equation 1.



**Fig. 1** Ambient occlusion as the cumulative distribution function.


#### 3.1 Ambient occlusion as CDF

The ratio of the volumes of the unoccluded points of the tangent sphere and the volume of the tangent sphere is the

geometric probability that a uniformly distributed point  $\mathbf{p}$  in the tangent sphere belongs to the unoccluded region, i.e. its scalar value  $v(\mathbf{p})$  is smaller than the scalar value  $v_{\mathbf{x}}$  of the shaded point  $\mathbf{x}$  or the isosurface. If the scalar field is defined by a voxel grid, then we have exact values at voxel centers, but have no information about the function  $v(\mathbf{p})$  in between the voxel centers. Note that the usual assumption on tri-linearly varying function is just an approximation. Instead, we can also assume that the scalar value  $v(\mathbf{p})$  is a random variable whose distribution is defined by the  $N$  voxel values  $v_1, \dots, v_N$  included in the tangent sphere. With these, the geometric probability of the occlusion is

$$O(\mathbf{x}) = P(v \leq v_{\mathbf{x}}) = \frac{\sum_{i=1}^N \varepsilon(v_{\mathbf{x}} - v_i)}{N},$$

where  $\varepsilon(x)$  is the step function. Note that  $P(v \leq v_{\mathbf{x}}) = F(v_{\mathbf{x}})$  is the cumulative probability distribution of random variable  $v$  for value  $v_{\mathbf{x}}$  (Figure 1).



**Fig. 2** The ambient occlusion of a voxel is approximated by computing the probability that the density of a voxel in the bounding box of the tangent sphere is greater than the density of the processed voxel.

Unfortunately, the application of this formula requires  $N$  fetches from the voxel array, which is time consuming unless  $N$  is small. However, selecting a small tangent sphere that contains just a few voxel centers degrades the quality of ambient occlusion and replaces nice soft shadows by smaller hard ones. On the other hand, when the set of voxels contained by the tangent sphere changes, there is an abrupt change in the ambient occlusion values, which results in stripe artifacts. So it is highly desirable to find a continuous approximation of the probability distribution, which is defined by just a few representative variables [2]. Keeping real-time rendering in mind even for dynamically evolving scalar fields, we should find these parameters to allow their computation with separable filtering, so the computation cost of a single representative variable is just  $\mathcal{O}(N^{1/3})$  instead of  $\mathcal{O}(N)$ . To make separable filtering possible, we extend the domain from the tangent sphere to its bounding box (see Figure 2).

### 3.2 Reconstruction of the cumulative distribution

In this approach we reconstruct the cumulative distribution from three quantities that can be obtained by separable filtering inside the bounding box of the tangent sphere: the minimum density value  $v_{\min}$ , the maximum density value  $v_{\max}$  and the mean of the density values  $E[v]$ .

The yet unknown cumulative distribution  $F(v_{\mathbf{x}})$  must be zero if  $v_{\mathbf{x}} \leq v_{\min}$ , equal to 1 if  $v_{\mathbf{x}} \geq v_{\max}$ , and non-decreasing in between. For notational simplicity, we introduce the following normalized parameter

$$t = \frac{v_{\mathbf{x}} - v_{\min}}{v_{\max} - v_{\min}} = \frac{v_{\mathbf{x}} - v_{\min}}{\Delta v},$$

where  $\Delta v = v_{\max} - v_{\min}$ .

Using the normalized parameter, the cumulative distribution must be zero if  $t \leq 0$ , and 1 if  $t \geq 1$ , and non-decreasing in interval  $[0, 1]$ . Considering all possible distribution of the density values, the cumulative distribution may be a step function at  $v_{\min}$  at one extreme case and another step function at  $v_{\max}$  in the other extreme case. Thus, we have to find a function that increases from 0 to 1 and has the flexibility to adapt to the two extreme cases. A very simple function meeting all these requirements is  $t^\beta$  where  $\beta$  is the parameter of data fitting (Figure 3). Thus the cumulative distribution is

$$F(v_{\mathbf{x}}) = t^\beta \quad \text{where} \quad v_{\mathbf{x}}(t) = t\Delta v + v_{\min}.$$

Let us consider the constraint on the mean:

$$E[v] = \int_{v_{\min}}^{v_{\max}} v_{\mathbf{x}} dF = \int_0^1 v_{\mathbf{x}}(t) \frac{dF}{dt} dt = \int_0^1 (t\Delta v + v_{\min}) \beta t^{\beta-1} dt = \frac{\beta \Delta v}{\beta + 1} + v_{\min}.$$

Solving this equation for  $\beta$ , we get:

$$\beta = \frac{E[v] - v_{\min}}{v_{\max} - E[v]}.$$

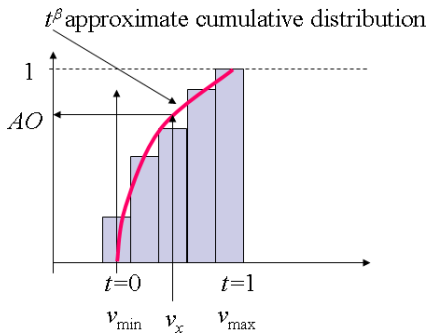


Fig. 3 Approximation of the cumulative distribution function with  $t^\beta$ .




Fig. 4 The Gaussian CDF with different parameters.

### 3.3 Gaussian cumulative distribution

Our second approach is to approximate the distribution by a Gaussian, and thus use the Gaussian CDF:

$$\begin{aligned} F(x; \mu, \sigma^2) &= \int_{-\infty}^x f(t; \mu, \sigma^2) dt \\ &= \Phi\left(\frac{x - \mu}{\sigma}\right) \\ &= \frac{1}{2} \left[ 1 + \operatorname{erf}\left(\frac{x - \mu}{\sigma\sqrt{2}}\right) \right], \quad x \in \mathbb{R}, \end{aligned} \quad (4)$$

where  $\mu = E[x]$ ,  $\sigma^2 = E[x^2] - E^2[x]$ , and  $\operatorname{erf}(x)$  is the *Gauss error function*, which can be approximated by

$$\operatorname{erf}(x) \approx \sqrt{1 - \exp\left(-x^2 \frac{4/\pi + ax^2}{1 + ax^2}\right)},$$

where


$$a = \frac{8(\pi - 3)}{3\pi(4 - \pi)} \approx 0.140012;$$

this approximation is better than  $4 \cdot 10^{-4}$  in relative precision [15].

In this case we compute by separable filtering the mean and the variance of the density values and use them in the previous equation to estimate the percentage of occluding voxels of a given voxel.

## 4 Results

The proposed methods have been implemented in CUDA and their performance has been measured on a system with an Intel Core 2 Quad Q9550 CPU and a Nvidia GeForce GTX 280 graphics card. To evaluate them we have used two different datasets, a CT-body of  $256 \times 256 \times 415$  voxels from the Visible Human project and the Chapel Hill CT-head of  $256 \times 256 \times 113$  voxels from the Volume Library.



**Fig. 6** AO maps for CT-head using different methods and radii. Rows: (i) Penner, (ii) reconstruction of the CDF, and (iii) Gaussian CDF. Columns: (a) radius of 5 voxels, and (b) radius of 10 voxels.

#### 4.1 Ambient occlusion maps

The proposed approaches have been compared with three different ambient occlusion strategies: the methods of Stewart [13], Ruiz [12], and Penner [10]. In Figure 5 we present the obtained results. From column (a) to (e), respectively, the maps have been obtained with the methods of Stewart, Ruiz, Penner, the proposed CDF reconstruction, and the Gaussian CDF. In the case of Stewart and Ruiz, the maps have been obtained considering 42 viewpoints. The other maps have been obtained taking a radius of 10 voxels. Observe that CDF reconstruction and Gaussian CDF methods generate lighter and more pleasant maps than the Penner method, and a smoother shading than the Stewart and Ruiz maps.

One of the parameters that can be modified in our approaches is the radius of the filter kernel, that is taken equal to the diameter of the tangent sphere. To illustrate the effect of this value on the final images we present in Figure 6

	Radius	Penner	CDF rec.	Gaussian CDF
CT-body	5	713	565	717
	10	958	734	965
CT-head	5	200	161	203
	10	268	205	270

**Table 1** Times in milliseconds to compute the AO maps for the different ambient occlusion methods and datasets and two radii.

the maps obtained with different methods and radii. From (i) to (iii), the maps are generated with the methods of Penner, CDF reconstruction and Gaussian CDF. Columns (a) and (b) are obtained with a radius of 5 and 10 voxels, respectively. It can be seen that the quality of the maps is improved when the radius is increased. Also, we can see that CDF reconstruction and Gaussian CDF have a greater quality than Penner's approach. Finally, we consider the ambient occlusion value as an ambient lighting  $AL$  term. In this case the color of a voxel  $\mathbf{x}$  is obtained as

$$C(\mathbf{x}) = AL(\mathbf{x}) = k_i O(\mathbf{x}) C_{TF}(\mathbf{x}), \quad (5)$$

where  $k_i$  is a constant factor that modulates the intensity of ambient occlusion  $O(\mathbf{x})$  and  $C_{TF}(\mathbf{x})$  is the pure color of the voxel as defined in the transfer function. In Figures 7 and 8 we illustrate the effect that can be obtained with the CT-body and the CT-head, respectively. From (a) to (d), respectively, we show the model without ambient occlusion, and applying the AO maps obtained with the methods of Penner, CDF reconstruction, and Gaussian CDF, in all cases with a radius of 10. As expected, by adding the ambient occlusion maps, the overall features of the volume models become more distinguishable. With the CDF reconstruction and the Gaussian CDF methods, context information is better captured, giving an improved depth perception over Penner's method, which generates a too dark final appearance. This is because, in the Chebychev estimation used by Penner's method, voxels with a value lower than the mean are assumed to be not occluded, and this creates a big contrast.

Table 1 shows the times in milliseconds to compute the volumetric AO maps for the CT-body and CT-head datasets, with radii 5 and 10, and for the three different methods. These times include all the filtering passes, memory transfers, and synchronizations.

#### 4.2 Opacity modulation

As Ruiz et al. [12] proposed, the gradient of the ambient occlusion indicates where the salient regions are, so it is worth modulating the opacity value by this gradient, which is estimated by a 4D regression filter [9]. The gradient calculation cannot be done using a separable filter, which increases the computational cost.

Instead of the gradient we propose the application of the variance of the intensity value, which also indicates where

the volume changes quickly. The variance can be obtained as the difference of the second moment and the square of the mean value and both of them can be obtained by separable filtering. Thus modified opacity  $\alpha'$  is:

$$\alpha' = \alpha \cdot \min\left(\frac{\sigma^2}{\sigma_{\max}^2}, 1\right) = \alpha \cdot \min\left(\frac{E[v^2] - E^2[v]}{\sigma_{\max}^2}, 1\right)$$

where  $\sigma_{\max}^2$  is a user defined constant.

In Figures 9 and 10 we present the opacity modulation effects that can be obtained using different methods. From column (a) to (e), respectively, we show the original model rendered with diffuse lighting, and modulated using the gradient of the AO map obtained with the methods of Penner, CDF reconstruction, and Gaussian CDF, and finally, with the intensity variance using a radius of 10 voxels and  $\sigma_{\max}^2 = 5$ . The result of this modulation is that the opacity of less occluded parts, such as muscles, is decreased, allowing to better see the inner structures. The time for the intensity variance computation with a radius of 10 voxels is 863 ms for the CT-body and 243 ms for the CT-head.


## 5 Conclusions

Volumetric ambient occlusion is a powerful technique that computes ambient occlusion by considering the portion of tangent sphere that is unoccluded. We have presented here new methods for the evaluation of the volumetric ambient occlusion for volumetric models. They produce high quality smooth shading and work at interactive rates thanks to the use of separable filtering to compute the needed parameters, and compare well with other ambient occlusion techniques for volumetric data. We have also shown how the gradient of the volumetric ambient occlusion nicely modulates the transfer function, and finally we have given a quick way to simulate this gradient.

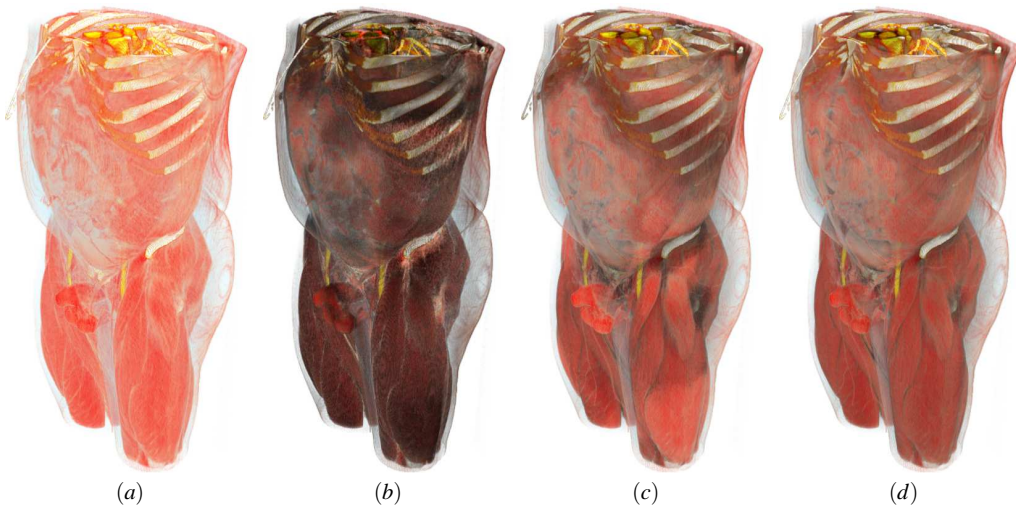
**Acknowledgements** This work was supported by TIN2007-68066-C04-01 and TIN2007-67982-C02 of the Ministry of Education and Science (Spanish Government) and by the 2009 SGR 643 from the Generalitat of Catalunya. The CT-head data set is courtesy of North Carolina Memorial Hospital, US (<http://www9.informatik.uni-erlangen.de/External/vollib/>). The CT-body data set is courtesy of the Visible Human Project, National Library of Medicine, US (<http://www.nlm.nih.gov/research/visible/>).

## References


- Desgranges, Engel: Fast ambient occlusion for direct volume rendering. US patent application 2007/0013696A1 (2007)
- Grun, H.: Approximate cumulative distribution function shadow mapping. In: W. Engel (ed.) *Shader X 7*. Course Technology (2008)
- 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*, pp. 1–8 (2007)
- Iones, A., Krupkin, A., Sbert, M., Zhukov, S.: Fast, realistic lighting for video games. *IEEE Comput. Graph. Appl.* **23**(3), 54–64 (2003)
- Jainek, W., Born, S., Bartz, D., Straßer, W., Fischer, J.: Illustrative Hybrid Visualization and Exploration of Anatomical and Functional Brain Data. pp. 855–862 (2008)
- Landis, H.: Production-ready global illumination. In: *Course 16 notes, SIGGRAPH 2002* (2002)
- Àlex Méndez-Feliu, Sbert, M., Neumann, L.: Obscurances for ray-tracing. In: *EUROGRAPHICS 2003 Poster Presentations*. Granada, Spain (2003)
- Mendez-Feliu, A., Sbert, M.: From obscurances to ambient occlusion: A survey. *The Visual Computer* **25**(2), 181–196 (2008)
- Neumann, L., Csébfalvi, B., König, A., Gröller, E.: Gradient estimation in volume data using 4d linear regression. *Comput. Graph. Forum* **19**(3), 351–358 (2000)
- Penner, E., Mitchell, R.: Isosurface ambient occlusion and soft shadows with filterable occlusion maps. In: *IEEE/EG Symposium on Volume and Point-Based Graphics*, pp. 57–64 (2008)
- Ropinski, T., Meyer-Spradow, J., Diepenbrock, S., Mensmann, J., Hinrichs, K.H.: Interactive volume rendering with dynamic ambient occlusion and color bleeding. *Computer Graphics Forum (Eurographics 2008)* **27**(2), 567–576 (2008)
- Ruiz, M., Boada, I., Viola, I., Bruckner, S., Feixas, M., Sbert, M.: Obscurance-based volume rendering framework. In: *Proceedings of IEEE/EG International Symposium on Volume and Point-Based Graphics*, pp. 113–120 (2008)
- Stewart, A.J.: Vicinity shading for enhanced perception of volumetric data. In: *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, pp. 355–362 (2003)
- Szirmay-Kalos, L., Umenhoffer, T., Toth, B., Szecsi, L., Sbert, M.: Volumetric ambient occlusion for real-time rendering and games. *Computer Graphics and Applications, IEEE* **30**(1), 70–79 (2010)
- Winitzki, S.: A handy approximation for the error function and its inverse. <http://homepages.physik.uni-muenchen.de/~winitzki/erf-approx.pdf> (2008)
- Wyman, C., Parker, S.G., Shirley, P., Hansen, C.D.: Interactive display of isosurfaces with global illumination. *IEEE Trans. Vis. Comput. Graph.* **12**(2), 186–196 (2006)
- Zhukov, S., Iones, A., Kronin, G.: An ambient light illumination model. In: *Proceedings of Eurographics Rendering Workshop '98*, pp. 45–56 (1998)



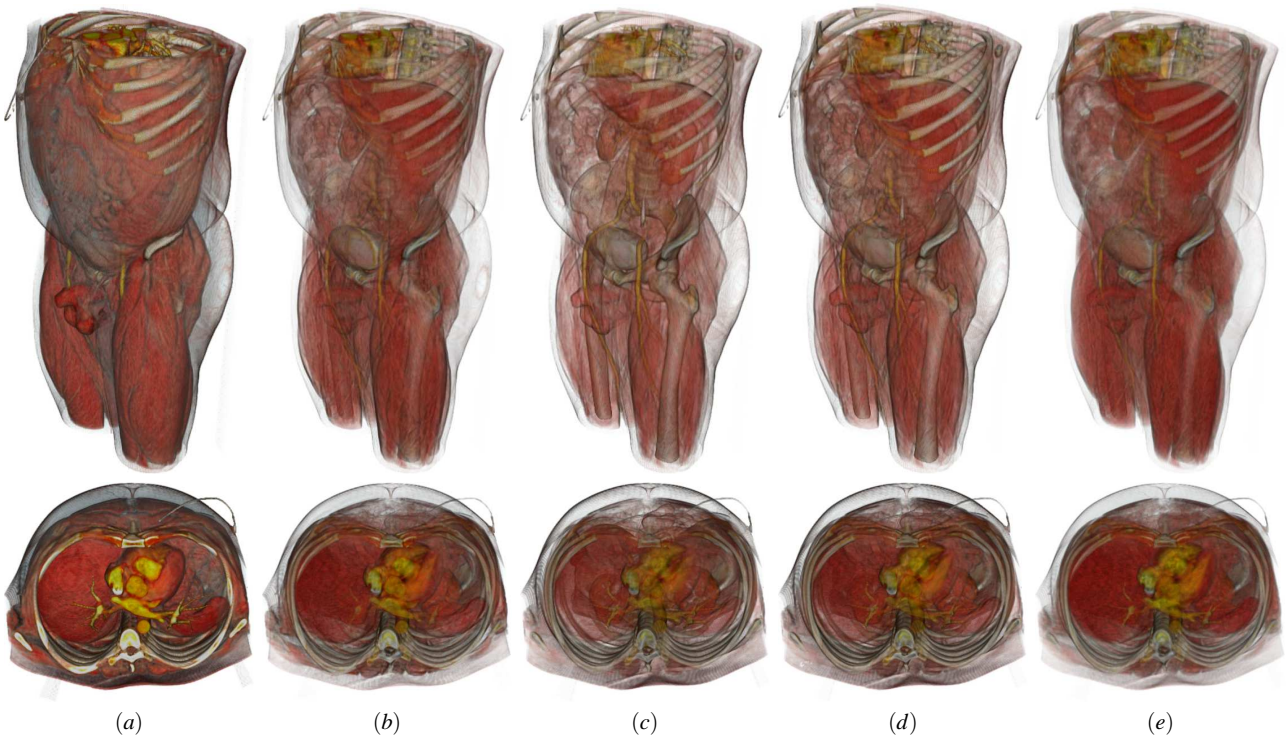
**Fig. 5** AO maps for CT-body using different methods: (a) Stewart, (b) Ruiz, (c) Penner, (d) reconstruction of the CDF, and (e) Gaussian CDF. (a) and (b) have been computed with 42 viewpoints, and (c), (d) and (e) with a radius of 10 voxels.




**Fig. 7** Ambient lighting for CT-body using different methods and a radius of 10 voxels: (a) original without ambient occlusion, (b) Penner's, (c) reconstruction of the CDF, and (d) Gaussian CDF.



**Fig. 8** Ambient lighting for CT-head using different methods and a radius of 10 voxels: (a) original without ambient occlusion, (b) Penner's, (c) reconstruction of the CDF, and (d) Gaussian CDF.



**Fig. 9** Opacity modulation of CT-body using different methods: (a) original rendering with diffuse lighting, (b) gradient of Penner's AO map, (c) gradient of CDF-reconstruction AO map, (d) gradient of Gaussian-CDF AO map, and (e) intensity variance with a radius of 10 voxels.



**Fig. 10** Opacity modulation of CT-head using different methods: (a) original rendering with diffuse lighting, (b) gradient of Penner's AO map, (c) gradient of CDF-reconstruction AO map, (d) gradient of Gaussian-CDF AO map, and (e) intensity variance.