Vision, Modeling, and Visualization (2012)
M. Goesele, T. Grosch, B. Preim, H. Theisel, and K. Toennies (Eds.)

Illumination-Driven Opacity Modulation for Expressive Volume Rendering

Balázs Csébfalvi[1], Balázs Tóth[1], Stefan Bruckner[2], Eduard Gröller[2]

[1]Budapest University of Technology and Economics, Hungary
[2]Vienna University of Technology, Austria

Abstract

Using classical volume visualization, typically a couple of isosurface layers are rendered semi-transparently to show the internal structures contained in the data. However, the opacity transfer function is often difficult to specify such that all the isosurfaces are of high contrast and sufficiently perceivable. In this paper, we propose a volume-rendering technique which ensures that the different layers contribute to fairly different regions of the image space. Since the overlapping between the effected regions is reduced, an outer translucent isosurface does not decrease significantly the contrast of a partially hidden inner isosurface. Therefore, the layers of the data become visually well separated. Traditional transfer functions assign color and opacity values to the voxels depending on the density and the gradient. In contrast, we assign also different illumination directions to different materials, and modulate the opacities view-dependently based on the surface normals and the directions of the light sources, which are fixed to the viewing angle. We will demonstrate that this model allows an expressive visualization of volumetric data.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture.

1. Introduction

The major goal of direct volume rendering is to represent the internal structures of the input 3D data by 2D images. This is a challenging task, since the volume is mapped to a one dimension lower representation, potentially leading to a loss of information. In order to avoid that important features get hidden, the isosurface layers of the data are rendered semi-transparently. If the opacities assigned to the different materials are relatively low then, theoretically, all the voxels contribute to the image. On the other hand, an isosurface of a low opacity is hardly perceivable due to its foggy appearance. The contrast can be increased by increasing the opacity, but then the partially hidden isosurfaces become less apparent. The different layers are difficult to distinguish especially if they contribute to the same regions of pixels. To remedy this problem, we propose a volume-rendering technique that well distributes the visual contributions of the isosurfaces in the image space. Therefore, the visual separation of the different structures becomes easier. The key idea is to illuminate the different layers from different directions and modulate their opacities depending on the illumination direction. To the best of our knowledge, this approach has not been considered for improving volume rendering so far.

2. Related Work

Due to the inherent complexity of volume data, resolving the problem of occlusion has been an important research topic in the area of volume visualization. A simple approach for removing unwanted occluders is clipping. Weiskopf et al. [WEE03] presented techniques for interactively applying arbitrary convex and concave clipping objects. Konrad-Verse et al. [KVPL04] proposed the use of a mesh which can be flexibly deformed by the user with an adjustable sphere of influence. Zhou et al. [ZDT04] applied distance-based opacity modulation to emphasize and deemphasize different regions. The work of Viola et al. [VKG05] introduced a method for automatically adapting the representation of contextual structures based on occlusion relationships resulting in visualizations similar to common cutaway views.

© The Eurographics Association 2012.

DOI: 10.2312/PE/VMV/VMV12/103-109

delivered by
EUROGRAPHICS DIGITAL LIBRARY
www.eg.org diglib.eg.org

Another way of resolving occlusion is to employ a sparser representation which selectively emphasizes features. The seminal work of Levoy [Lev88] introduced the idea of modulating the opacity at a sample position using the magnitude of the local gradient. As homogeneous regions are suppressed, this effectively enhances boundaries in the volume. Rheingans and Ebert [RE01] presented several illustrative techniques which enhance features and add depth and orientation cues. Csébfalvi et al. [CMH*01] enhanced object contours based on the magnitude of local gradients as well as on the angle between the viewing direction and the gradient vector using depth-shaded maximum intensity projection. The concept of two-level volume rendering, proposed by Hauser et al. [HMBG01], allows focus+context visualization of volume data. This approach combines different rendering styles, such as direct volume rendering and maximum intensity projection.

A way for defining transfer functions based on occlusion information was introduced by Correa and Ma [CM09]. The occlusion spectrum enables volume classification based on the ambient occlusion of voxels. By assigning different opacities to certain degrees of occlusion, hidden structures can be uncovered. In further work [CM11], they also presented the use of visibility histograms which allow semi-automatic generation of transfer functions which maximize the visibility of important structures.

Our work makes use of illumination information to enable the uncluttered simultaneous depiction of multiple structures of interest. Lum and Ma [LM04] proposed to assign colors and opacities as well as parameters of the illumination model through a transfer function lookup. They applied a two-dimensional transfer function to emphasize material boundaries using illumination. Their work only allows the specification of varying ambient, diffuse, and specular reflection properties for different structures in the volume. Bruckner and Gröller [BG05] parameterized the shading model to allow simple generation of illustrative effects such as metallic shading. They also introduced style transfer functions [BG07] which combine color and shading in a single transfer function defined by multiple sphere maps.

Context-preserving volume rendering, as introduced by Bruckner et al. [BGKG06], employed the idea of using illumination as an input to the opacity function. While our work is based on a similar notion, we use selective inconsistent illumination in order to maximize the visibility of different features in the volume. One inspiration for our approach is the work of Lee at al. [LHV06]. Their system optimizes the placement of local light sources to enhance the curvature of a polygonal mesh. In this paper, we demonstrate how inconsistent lighting can be used to enhance the visibility of multiple structures in a volume data set.

3. Illumination-Driven Opacity Modulation

Although our new volume-rendering technique is not physically plausible, for the sake of clarity, we still explain it through a fictional physical model. Afterwards, we discuss how to use this model with an appropriate lighting design for expressive volume visualization.

3.1. Fictional Physical Model

We assume that n isosurfaces need to be rendered. We apply the same number of light sources emitting invisible light of different frequencies f_i, where $i \in \{1, 2, ..., n\}$. None of the isosurfaces attenuate the invisible light rays; thus, the incoming invisible light is completely transmitted without changing its direction, frequency, or intensity. However, each isosurface is sensitive to exactly one frequency. More concretely, if an isosurface s_i is hit by a light ray of frequency f_i then the incoming invisible light is not just completely transmitted, but additionally reflected and its frequency is changed such that it becomes visible. The reflection direction depends on the Bidirectional Reflectance Distribution Function (BRDF) assigned to s_i. Consequently, the shaded color of the intersection point is calculated by evaluating a local shading model according to the BRDF taking the surface normal, the direction of the incoming light, and the viewing direction into account. Although the isosurfaces fully transmit the incoming invisible light, they do attenuate the reflected visible light. Therefore, we assign opacity values to each intersection point depending on how much the given point is illuminated by the corresponding light source. Without a loss of generality, assume that the isosurfaces are diffuse. In this case, based on the Lambertian shading model, the illumination intensity is calculated as $\mathbf{max}(\mathbf{N} \cdot \mathbf{L}, 0)$, where $\mathbf{N}$ is the surface normal and $\mathbf{L}$ is the direction of the light source. The opacity values assigned to the isosurfaces are then modulated by this illumination intensity.

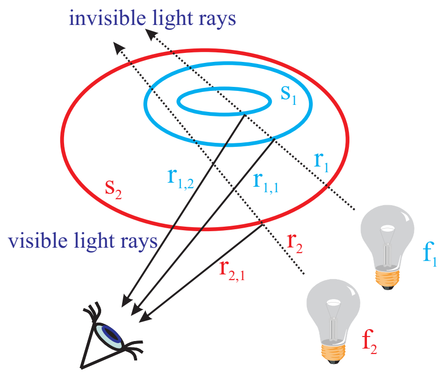

Figure 1: *Illustration of our fictional optical model.*

© The Eurographics Association 2012.

Figure 1 shows the illustration of our model for two light sources ($n = 2$). Invisible light rays r_1 and r_2 are of frequencies f_1 and f_2, respectively. Both of them are transmitted by the isosurfaces s_1 and s_2. As isosurface s_1 is sensitive to frequency f_1, it reflects r_1. The reflected rays originated from the first and second intersection points are denoted by $r_{1,1}$ and $r_{1,2}$, respectively. Since these reflected rays are visible, they are attenuated by isosurface s_2. Because of the intersection with the outer layer of s_1, reflected ray $r_{1,2}$ is attenuated by isosurface s_1 as well. The attenuation factors depend on the normal vectors at the corresponding intersection points and the direction of the light source the given isosurface interacts with. Isosurface s_2 is sensitive only to frequency f_2. Therefore, it reflects only the invisible light ray r_2. The reflected ray $r_{2,1}$ is visible and reaches the eye position without attenuation.

For an arbitrary number of isosurfaces, compositing of the visible rays is implemented by the algorithm shown in Listing 1.

```
COLOR RayCompositing(VOLUME volume, RAY ray)
{
  COLOR output = BLACK;
  VECTOR position = ray.origin;
  REAL transparency = 1.0;
  REAL density = 0.0;

  // front-to-back compositing along the ray
  while(IsInBoundingBox(position))
  {
    REAL prevDensity = density;
    position += ray.direction;
    REAL density = volume.Resample(position);

    COLOR color;
    REAL opacity;

    // for each isosurface
    for(int i = 0; i < n; i++)
    {
      REAL threshold = surface[i].threshold;

      // the ith isosurface is intersected
      if((prevDensity - threshold) *
         (density - threshold) < 0)
      {
        color = surface[i].color;
        VECTOR normal = volume.NormalResample(position);
        VECTOR lightDir = lightSource[i].direction;
        REAL lighting = max(dot(normal, lightDir), 0);
        color *= lighting;

        // illumination-driven opacity modulation
        opacity = lighting * surface[i].opacity;
      }
    }

    color *= transparency * opacity;
    transparency *= 1.0 - opacity;
    oputput += color;
  }

  return output;
}
```

Listing 1: *Ray compositing with illumination-driven opacity modulation.*

The outer `while` loop goes through the samples along the ray in front-to-back order. The densities of the current and the previous samples are represented by variables `density` and `prevDensity`, respectively. An intersection point with an isosurface is detected if `density` is greater than the isosurface threshold and `prevDensity` is lower, or vice versa. The inner `for` loop goes through all the isosurfaces and checks the potential intersection points. At each intersection point, a normal is evaluated by gradient estimation, and a lighting coefficient is calculated depending on the directions of the normal and the corresponding light source. The color of an intersection point is the diffuse color assigned to the given isosurface modulated by the lighting coefficient. The opacity is calculated similarly, i.e., the opacity value assigned to the given isosurface is modulated by the lighting coefficient as well. After having the color and opacity values defined for each intersection point, the standard front-to-back compositing is performed.

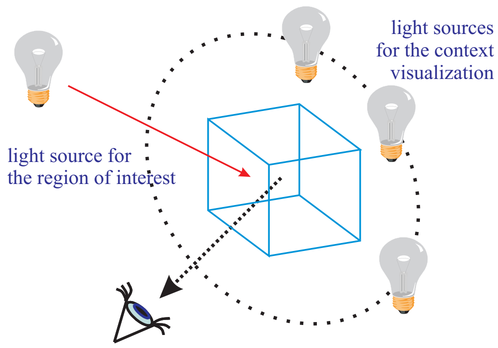

Figure 3: *Arrangement of the light sources around the volume.*

3.2. Illumination Design

The goal of the illumination design is to enhance the region of interest and to preserve the context information at the same time. To render the context, we propose to locate the corresponding light sources along a circle, which is perpendicular to the major viewing direction and has its center in the middle of the volume (see Figure 3). This arrangement of light sources combined with our illumination-driven opacity modulation guarantees that each outer isosurface appears with high contrast and does not significantly suppress the inner isosurfaces. This is demonstrated by Figure 2 that shows three concentric spherical isosurfaces. Using traditional volume rendering, the different layers are of low contrast and their 3D nature is hardly perceivable. Practically, only the inner most red layer seems to be a well-defined spherical surface, while the two outer layers look like fog around the red sphere rather than sharp boundary surfaces. In contrast, applying illumination-driven opacity modulation, the 3D structure of each isosurface is clearly comprehensible. The red and green layers are illuminated from a vertical direction, whereas the blue layer is illuminated from a horizontal direction. Therefore, the different isosurfaces contribute to fairly different regions of the image space. This

© The Eurographics Association 2012.

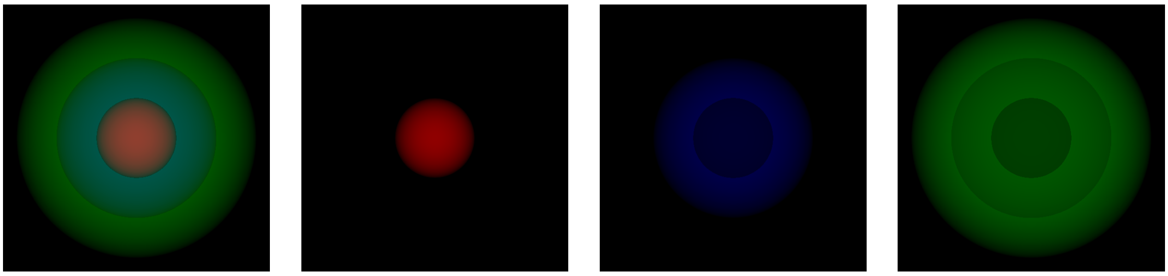

Traditional volume rendering.

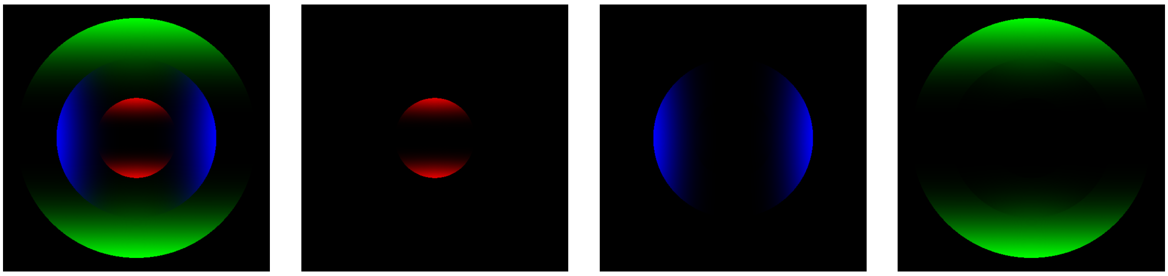

Illumination-driven opacity modulation.

Figure 2: *Comparison of traditional volume rendering to illumination-driven opacity modulation. The contributions of the different isosurfaces are also shown separately. Note that our illumination-driven opacity modulation approach significantly enhances the contrast of each isosurface.*

makes the visual separation of the layers much easier than in case of traditional volume rendering. There the visual separation is supposed to be guaranteed only through the color and opacity information. Even though the red, green, and blue color channels are assigned separately to the three isosurfaces, none of them can fully exploit the range of its own color channel because the constant opacity modulation drastically limits the highest possible brightness.

The isosurfaces belonging to the region of interest should not necessarily be lit perpendicularly to the viewing direction. In this case, the illumination direction can be interactively controlled on the fly. In Section 5, we will demonstrate how this additional degree of freedom can be utilized for volumetric data exploration. In fact, the light source assigned to the region of interest is used as a "magic lamp", which can enhance different subsurfaces of the internal structures depending on the illumination direction.

4. Experimental Results

We tested our volume-rendering model on real-world CT data. We intentionally chose test data sets, in which the density thresholds of two important isosurfaces are relatively close to each other. Therefore, their visual separation is a difficult task. One of the test data sets is a CT scan of a human hand, while the other one is the CT scan of a human head. In both data sets, the blood vessels contain contrast agent. Since the density of the contrast agent is nearly the same as that of the bone, it is not easy to separate the blood vessels from the bone structures. The upper two images in Figure 4 show the results generated by classical volume rendering. It is clearly apparent that the blood vessels can hardly be distinguished from the bones based on only the color information assigned according to the densities. However, it is still worthwhile to show the lower-density membrane around the bones and the blood vessels with red color. Note that, using classical alpha-blending with constant opacities, this red membrane is visible only around the thinner vessels, so the meaning of the color information becomes incoherent. Instead, we propose to illuminate the two layers from different directions (see the lower two images in Figure 4). To visualize the hand data set, we illuminate the white layer from the viewing direction and the red layer from a horizontal direction. The yellowish isosurface of the skin, is also illuminated horizontally. Here the interpretation of the red regions is quite clear, as they visually well represent the lower-density membranes. It is also important to mention that the horizontally illuminated skin layer, which represents the context information, does not significantly hide the internal structures (in other words, the region of interest) but its shape is still perceivable. To visualize the head data set, we used almost the same illumination settings as for the hand, but the skull is illuminated vertically and not from the viewing direction, to avoid that it hides the internal blood vessels. It is clearly visible that,

© The Eurographics Association 2012.

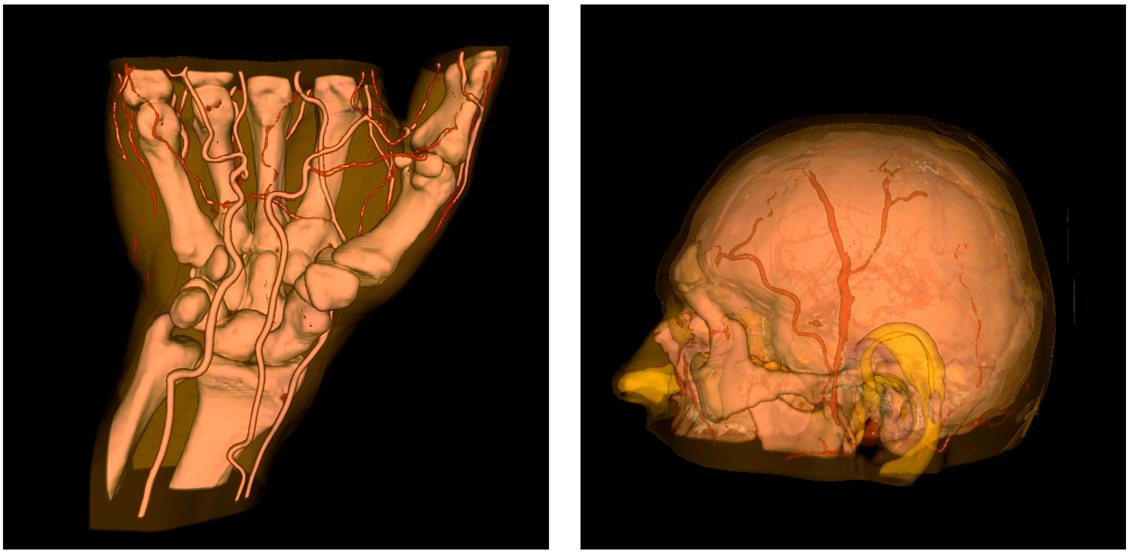

Traditional volume rendering.

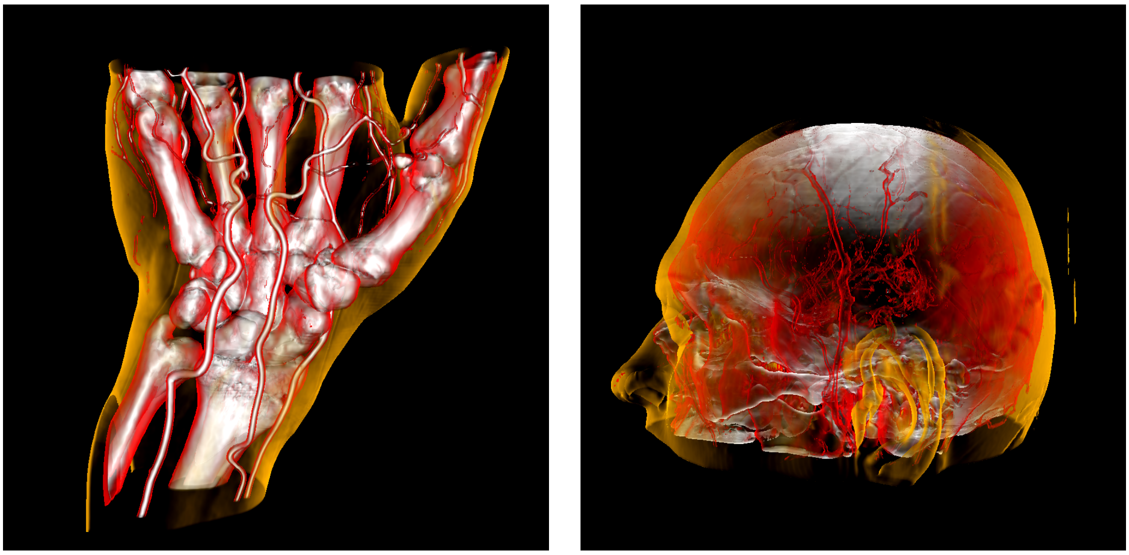

Illumination-driven opacity modulation.

Figure 4: *Comparison of traditional volume rendering to illumination-driven opacity modulation on real-world CT data. Note that, unlike traditional volume rendering, our illumination-driven opacity modulation ensures high contrast for each isosurface.*

compared to classical volume rendering, our visualization model results in images richer in detail. In Figure 6, the contributions of the three different layers are shown separately. Note that, using illumination-driven opacity modulation, the contrast of each isosurface is very well preserved for both test data sets. Such a high contrast is not at all provided by classical volume rendering, which significantly reduces the range of colors because of the constant opacity modulation.

5. Interactive Illumination Control

Due to the efficient GPU implementation, the illumination directions can be interactively modified (rendering images of resolution 512×512, we measured frame rates of 8-10 fps on an NVIDIA GeForce GTX 480 graphics card). Especially the lighting of the region of interest is exciting to modify. This is illustrated in Figure 5, where the skin and the blood vessels are lit horizontally, while the illumination of the skull is varying between the vertical and viewing direc-

© The Eurographics Association 2012.

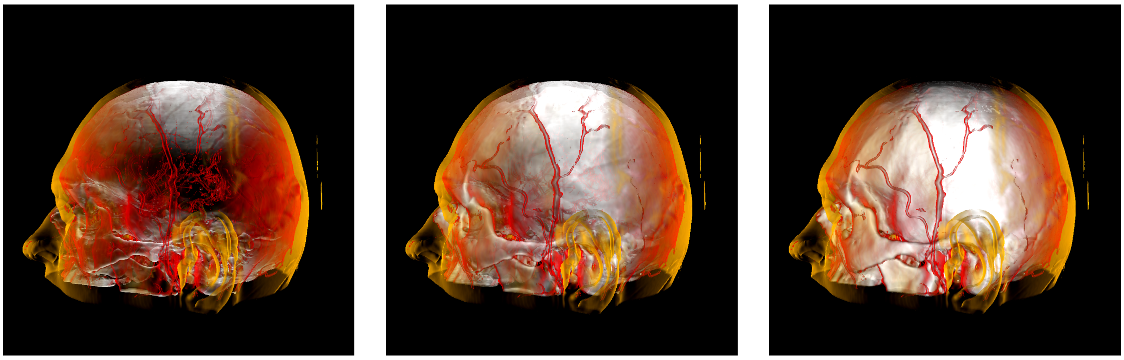

Figure 5: *Illustration of the interactive illumination control. The illumination of the skin and the blood vessels is fixed, while the illumination of the skull is varying between the vertical and viewing directions.*

tions. Note that, changing the directions of the light sources, the different features of the data can be enhanced or suppressed in a flexible way. We think that this kind of volume exploration is more intuitive than a simple modification of the transfer function parameters having the illumination directions fixed. Although we have not done a thorough user study yet, the physicians we consulted with found this interaction scheme useful.

6. Conclusion

In this paper, we have introduced illumination-driven opacity modulation for expressive volume rendering. We have shown that lighting the different isosurfaces from different directions guarantees that the layers are visually well separated, as they contribute to fairly different regions in the generated image. Moreover, our method well preserves the contrast of all the isosurfaces, so their visual interpretation becomes easier. Last but not least, using an efficient GPU implementation, the illumination directions can be interactively controlled, which is especially useful for enhancing the region of interest.

Acknowledgements

This work has been supported by projects TÁMOP-4.2.2.B-10/1–2010-0009, OTKA K-101527, OTKA K-104476, and the ViMaL project (FWF, no. P21695).

References

[BG05] BRUCKNER S., GRÖLLER M. E.: VolumeShop: An interactive system for direct volume illustration. In *Proceedings of IEEE Visualization 2005* (2005), pp. 671–678. 2

[BG07] BRUCKNER S., GRÖLLER M. E.: Style transfer functions for illustrative volume rendering. *Computer Graphics Forum 26*, 3 (2007), 715–724. 2

[BGKG06] BRUCKNER S., GRIMM S., KANITSAR A., GRÖLLER M. E.: Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization and Computer Graphics 12*, 6 (2006), 1559–1569. 2

[CM09] CORREA C. D., MA K.-L.: The occlusion spectrum for volume classification and visualization. *IEEE Transactions on Visualization and Computer Graphics 15* (2009), 1465–1472. 2

[CM11] CORREA C. D., MA K.-L.: Visibility histograms and visibility-driven transfer functions. *IEEE Transactions on Visualization and Computer Graphics 17*, 2 (2011), 192–204. 2

[CMH*01] CSÉBFALVI B., MROZ L., HAUSER H., KÖNIG A., GRÖLLER M. E.: Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum 20*, 3 (2001), 452–460. 2

[HMBG01] HAUSER H., MROZ L., BISCHI G. I., GRÖLLER M. E.: Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 242–252. 2

[KVPL04] KONRAD-VERSE O., PREIM B., LITTMANN A.: Virtual resection with a deformable cutting plane. In *Proceedings of Simulation und Visualisierung 2004* (2004), pp. 203–214. 1

[Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications 8*, 3 (1988), 29–37. 2

[LHV06] LEE C. H., HAO X., VARSHNEY A.: Geometry-dependent lighting. *IEEE Transactions on Visualization and Computer Graphics 12*, 2 (2006), 197–207. 2

[LM04] LUM E. B., MA K.-L.: Lighting transfer functions using gradient aligned sampling. In *Proceedings of IEEE Visualization 2004* (2004), pp. 289–296. 2

[RE01] RHEINGANS P., EBERT D. S.: Volume illustration: Non-photorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 253–264. 2

[VKG05] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven feature enhancement in volume visualization. *IEEE Transactions on Visualization and Computer Graphics 11*, 4 (2005), 408–418. 1

[WEE03] WEISKOPF D., ENGEL K., ERTL T.: Interactive clipping techniques for texture-based volume visualization and volume shading. *Visualization and Computer Graphics, IEEE Transactions on 9*, 3 (2003), 298–312. 1

[ZDT04] ZHOU J., DÖRING A., TÖNNIES K. D.: Distance based enhancement for focal region based volume rendering. In *Proceedings of Bildverarbeitung für die Medizin 2004* (2004), pp. 199–203. 1

© The Eurographics Association 2012.

Traditional volume rendering.

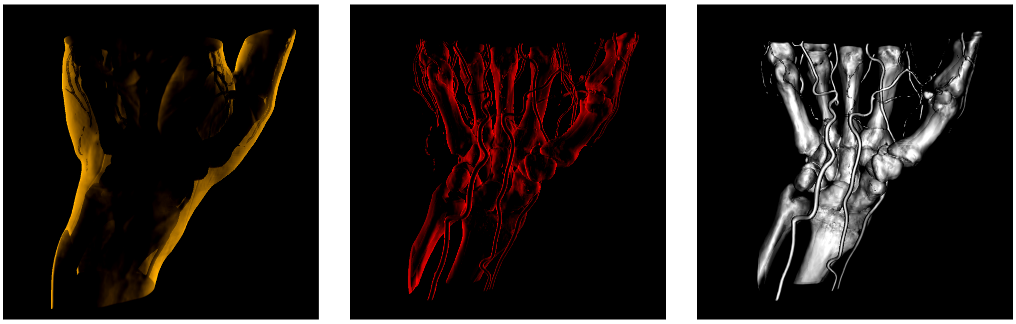

Illumination-driven opacity modulation.

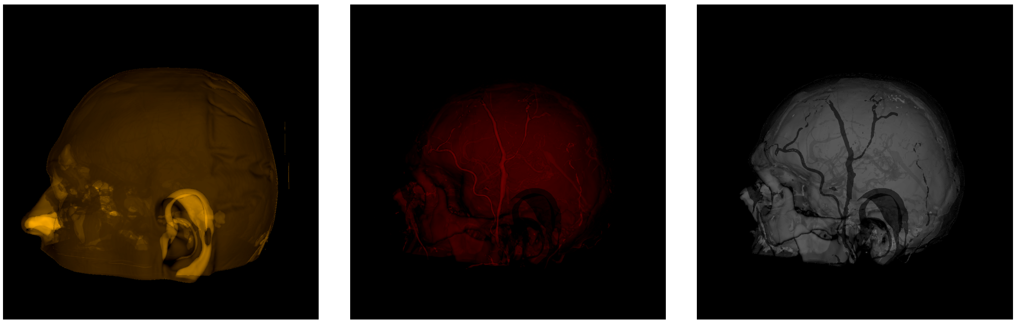

Traditional volume rendering.

Illumination-driven opacity modulation.

Figure 6: *Contributions of the different isosurfaces using traditional volume rendering and illumination-driven opacity modulation. Note that, in case of illumination-driven opacity modulation, the three different layers contribute to fairly different regions of the image space. Therefore, their visual separation becomes easier than in case of traditional volume rendering. Furthermore, our model well maintains the contrast of each isosurface.*

© The Eurographics Association 2012.