

A Perceptual Framework for Comparisons of Direct Volume Rendered Images

Hon-Cheng Wong¹, Huamin Qu², Un-Hong Wong¹,
Zesheng Tang¹, and Klaus Mueller³

¹ Faculty of Information Technology,

Macau University of Science and Technology, Macao, China

² Department of Computer Science and Engineering,

Hong Kong University of Science and Technology, Hong Kong, China

³ Department of Computer Science,

Stony Brook University, Stony Brook, USA

hcwong@ieee.org, huamin@cs.ust.hk, uhwong@computer.org,

ztang@must.edu.mo, mueller@cs.sunysb.edu

Abstract. Direct volume rendering (DVR) has been widely used by physicians, scientists, and engineers in many applications. There are various DVR algorithms and the images generated by these algorithms are somewhat different. Because these direct volume rendered images will be perceived by human beings, it is important to evaluate their quality based on human perception. One of the key perceptual factors is that whether and how the visible differences between two images will be observed by users. In this paper we propose a perceptual framework, which is based on the Visible Differences Predictor (VDP), for comparing the direct volume rendered images generated with different algorithms or the same algorithm with different specifications such as shading method, gradient estimation scheme, and sampling rate. Our framework consists of a volume rendering engine and a VDP component. The experimental results on some real volume data show that the visible differences between two direct volume rendered images can be measured quantitatively with our framework. Our method can help users choose suitable DVR algorithms and specifications for their applications from a perceptual perspective and steer the visualization process.

1 Introduction

Direct volume rendering (DVR) is a widely used technique in visualization, which directly renders 3D volume data into 2D images without generating any intermediate geometric primitives. There are many DVR algorithms developed in the past two decades, including ray-casting [1], splatting [2], shear-warp [3], 2D texture slicing [4], 3D texture slicing [5], and GPU-based volume rendering [4] [6] [7]. A recent survey of DVR algorithms can be found in [8].

It is well known that direct volume rendered images generated by different DVR methods are somewhat different and some algorithms can generate images with better quality than others. Therefore, there is a need to compare the quality of direct volume rendered images generated by different methods and specifications. Fortunately,

there are more and more works reporting the comparisons of volume rendering algorithms [9] [10] and volume rendered images [11].

These direct volume rendered images will be perceived by the human beings. Therefore, it is important to quantitatively evaluate DVR images based on human perception. One of the key perceptual factors is that whether the visible differences between two images will be observed. However, research concerning this factor is scant. In this paper we propose a perceptual framework, which is based on Daly's Visible Differences Predictor [12], for comparative study of direct volume rendered images. It will be used to predict the visible differences of the direct volume rendered images generated with different algorithms and the same algorithm with different specifications.

The remainder of this paper is organized as follows: We first introduce related work in Section 2. We then describe our framework and review VDP in Section 3. Next, we compare the direct volume rendered images by using our framework in Section 4. Finally, we conclude our work and discuss future research directions in Section 5.

2 Related Work

Comparative evaluation in DVR algorithms: Methods for comparing DVR algorithms can be categorized into three classes:

1) Image level methods [13] [9]: They usually compare DVR images side-by-side using various methods such as difference image, mean square errors (MSE), and root mean square error (RMSE); 2) Data level methods [10]: They use raw data and intermediate information obtained during the rendering process for comparison; 3) Analytical methods to calculate the error bounds [14]: They analyze the errors in gradient calculations, normal estimation schemes, and filtering and reconstruction operations. In addition to these comparative methods, Mei'sner et al. [15] performed a practical evaluation of popular DVR algorithms in terms of rendering performance on real-life data sets.

Perception issues in computer graphics: Considerable concern has arisen over the perception in graphics research in recent years, especially in the area of global illumination. There has been much work on perceptually-based rendering proposed. Most of them focus on two tasks: 1) To establish stopping criteria for high quality rendering systems by developing perceptual metrics [16] and 2) To optimally manage resource allocation for efficient rendering algorithms by using perceptual metrics [17] [18]. In addition, Rushmeier et al. [19] proposed some metrics for comparing real and synthetic images.

Perception issues in visualization: There is a growing number of research on using perception for visualization. For example, Lu et al. [20] utilized several feature enhancement techniques to create effective and interactive visualizations of scientific and medical data sets. Ebert [21], Interrante [22], and Chalmers and Cater [23] recently give excellent surveys on perception issues in visualization. Zhou et al. [24] presented a study of image comparison metrics for quantifying the magnitude of difference between a visualization of a computer simulation and a photographic image captured from an experiment.

To the best of our knowledge, there has been little work which applies perception in comparative evaluation of DVR. Inspired by the work of Myszkowski [18] which uses

VDP for global illumination problems, we have developed a framework based on VDP to compare the direct volume rendered images.

3 Our Framework and Visible Differences Predictor (VDP)

The block diagram of our framework is shown in Figure 1. Our framework contains a volume rendering engine which generates the images by using different DVR algorithms supported in the engine. Two direct volume rendered images are produced with the user specified settings. Then they will be sent to the VDP, which compares these images and gives a differences map (VDP responses) as output.

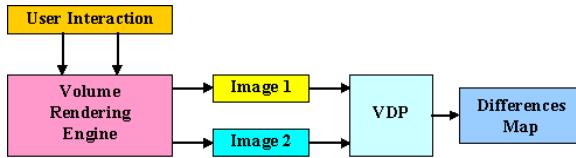


Fig. 1. Block diagram of our framework

There are many metrics based on Human Visual System (HVS). Two most popular ones are Daly's Visible Differences Predictor (VDP) [12] and Sarnoff Visual Discrimination Model (VDM) [25]. Both metrics were shown to perform equally well on average [26]. We chose VDP in our framework because of its modularity and extensibility. The block diagram of VDP [18] is shown in Figure 2. VDP receives as input a

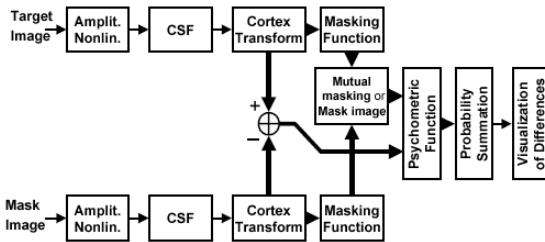


Fig. 2. Block diagram of the Visible Differences Predictor [18]

pair of images (target image and mask image), and outputs a map of probability values, which indicates how the differences between those images are perceived [12] [18]. These two input images are first processed by the *amplitude nonlinearity*, which simulates the adaptation of HVS to local luminance. Then the resulting images are converted into frequency domain using FFT. After that, *contrast sensitivity function (CSF)*, which simulates the variations in visual sensitivity of HVS, is performed on the frequency signals. These images are then converted to spatial frequency and orientation channels using a pyramid-style *cortex transform*. *Masking function* which is used to increase the

threshold of detectability is applied to these images and the minimal threshold elevation value for corresponding channels and pixels are taken by *mutual masking/mask image*. A *psychometric function* for predicting the probability of perceived differences is applied to these images and finally the predicted probability is visualized.

As VDP is a general purpose predictor of the differences between images, it can be used to evaluate pairs of images for a wide range of applications. Although VDP does not support chromatic channels in input images, in DVR applications many important insights such as depth cues can be well captured in an achromatic images. Thus we embedded VDP in our framework for comparing the direct volume rendered images. Since the HVS is more sensitive to the differences of contrast and less sensitive to the differences between colors, we convert color direct volume rendered images generated by the volume rendering engine to gray-scale images before sending them to the VDP.

All comparisons were performed on a 2.4GHz AMD Opteron 280 processor with 6GB main memory, and an NVIDIA Quadro FX 4500 graphics card with 512MB video memory. The resulting VDP responses (differences map in our framework) are represented as a color map which is blended with the original grey-scale target image. Color is added to each pixel in this target image to indicate its difference detection probability values. The probability values greater than 0.75, which is the standard threshold value for discrimination tasks [27], are set to red pixels. In the rest of the paper, we usually provide results in a set of three figures. The first two figures are generated using different algorithms or settings and the third one shows their differences map, which is encoded in the same color scale as in Figure 3 (d). The background pixels (black pixels) are not included in the calculation of percentage of red pixels in the VDP result.

4 Comparisons of Direct Volume Rendered Images

In this section, we use our framework to measure the perceptible differences in the direct volume rendered images generated with different algorithms or the same algorithm with different specifications. For DVR algorithms we select two most popular methods: GPU-based ray-casting [6] and 3D texture slicing [5]. And several specifications including shading, gradient estimation scheme, and sampling rate are chosen for comparisons. We limit our case studies to static, regular or rectilinear, scalar volume data only. The size of all images is 512×512 . All algorithm-independent parameters such as viewing, transfer functions, and optical model, are kept constant in each image comparison set in order to have a fair comparison. Experimental results will be discussed at the end of this section.

4.1 GPU-Based Ray-Casting Versus 3D Texture Slicing

Two data sets are used here: 256^3 CT human head and 256^3 MRI human head. Figure 3 and Figure 4 compare the direct volume rendered images generated with GPU-based ray-casting and 3D texture slicing for these two data sets respectively. From the VDP responses shown in Figure 3 (c), we can see that there are some noticeable areas of red pixels, which indicate the differences between Figure 3 (a) and (b) in these regions are quite noticeable (probability $> 75\%$) to human beings. And there are only some light green pixels inside the regions of red pixels. The percentage of red pixels is 11.89% in

the VDP result. In Figure 4 (c), we can see some pieces of red pixels and many small pieces of light green pixels on the surface of the MRI head. The percentage of red pixels is 28.27% in it. In both Figure 3 (c) and Figure 4 (c), the red pixels are distributed in the transparent regions. The VDP predictions (Figure 3 (c) and Figure 4 (c)) coincide with the human perception of the visual results shown in Figure 3 (a) and (b), and Figure 4 (a) and (b).

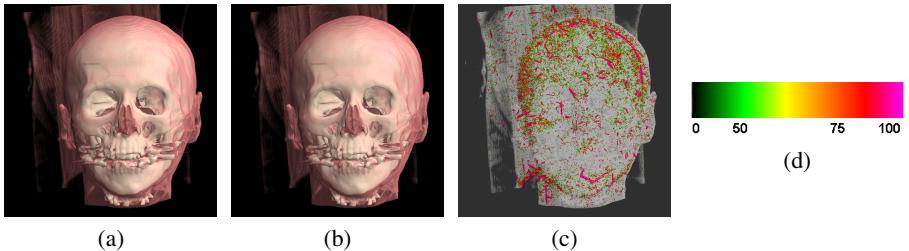


Fig. 3. Comparison of GPU-based ray-casting and 3D texture slicing (256^3 CT human head): (a) GPU-based ray-casting; (b) 3D texture slicing; (c) VDP result (Red pixels: 11.89%); (d) Color scales for encoding the probabilities (%) in (c)

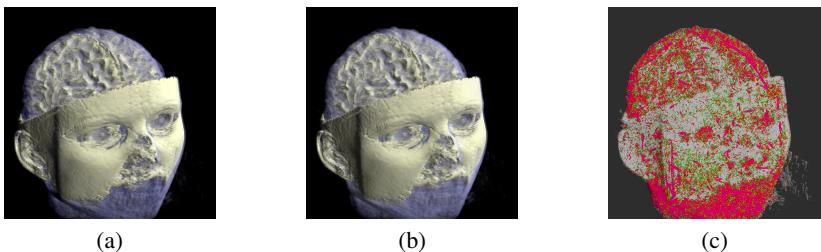


Fig. 4. Comparison of GPU-based ray-casting and 3D texture slicing (256^3 MRI human head): (a) GPU-based ray-casting; (b) 3D texture slicing; (c) VDP result (Red pixels: 28.27%)

4.2 Shading

We are interested in the visual differences between the rendering results in the following two scenarios. First, in pre-shaded DVR, the shading model at the grid samples is evaluated first and then the illumination is interpolated. In contrast, the normal is interpolated first and then the shading model for each reconstructed sample is evaluated in post-shaded DVR. Second, separate color interpolation was used in [1], which interpolates voxel colors and opacities separately before computing the product of them. Wittenbrink et al. [28] pointed out that it is more correct to multiply color and opacity beforehand at each voxel and then interpolate the product. We compare the pre-shaded and post-shaded DVR images, as well as the opacity-weighted color interpolated [28] and separate color interpolated DVR images using the GPU-based ray-caster with a 256^3 engine data set.

Figure 5 shows the comparison of pre-shaded and post-shaded DVR. From Figure 5 (c) we can know that the rendering results between pre-shaded (Figure 5 (a)) and post-shaded (Figure 5 (b)) DVR images are quite different. The percentage of red pixels reaches 54.79%, which means that such differences in these two images are very noticeable to human beings.

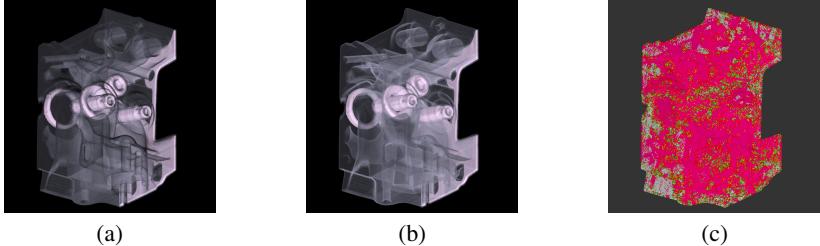


Fig. 5. Comparison of pre-shaded DVR and post-shaded DVR (256^3 engine): (a) Pre-shaded DVR; (b) Post-shaded DVR; (c) VDP result (Red pixels: 54.79%)

Figure 6 shows the comparison of opacity-weighted color interpolated and separate color interpolated DVRs. The percentage of red pixels is 64.00%. And the distribution of the most noticeable differences between Figure 6 (a) and (b) can be easily found in Figure 6 (c).

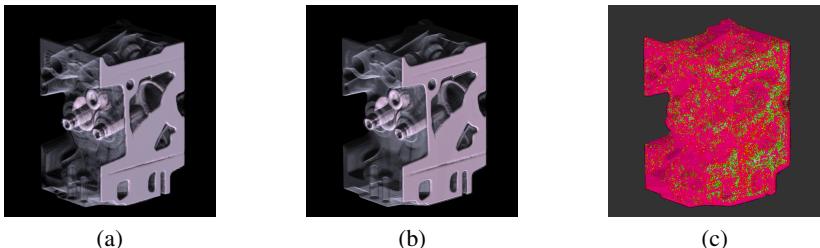


Fig. 6. Comparison of opacity-weighted color interpolated and separate color interpolated DVRs (256^3 engine): (a) Opacity-weighted color interpolated DVR; (b) Separate color interpolated DVR; (c) VDP result (Red pixels: 64.00%)

4.3 Gradient Estimation Scheme

In DVR, different gradient estimation scheme expresses the choice of normal computation from the volume data. They may significantly affect the shading and appearance of the rendering results. Two schemes are compared in Figure 7: central difference operator [1] which computes gradients at data values in the x , y , z direction and then uses the gradients at the eight nearest surrounding data locations to interpolate the gradient vectors for locations other than at data locations; intermediate difference operator [29] which computes gradient vectors situated between data locations using differences in data values at the immediate neighbors. Figure 7 (c) shows that the differences between the images are very obvious. The percentage of red pixels is 50.62%.

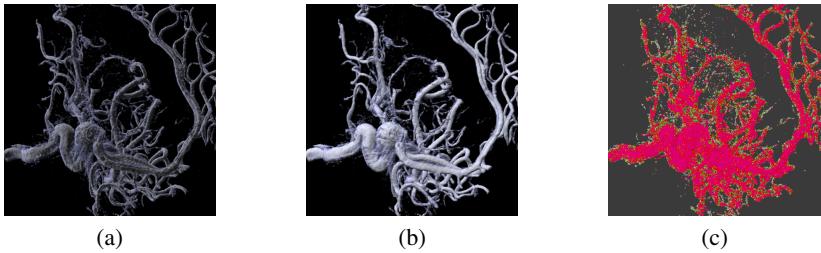


Fig. 7. Comparison of intermediate and central difference operators in DVR (256^3 aneurism): (a) With intermediate difference operator; (b) With central difference operators; (c) VDP result (Red pixels: 50.62%)

4.4 Sampling Rate

Direct volume rendered images generated by a GPU-based ray-caster with different sampling rates are compared in this section. The motivation for performing this kind of comparisons is that we want to reduce the rendering time of the data sets by downgrading their sampling rates without sacrificing the image quality too much. A 256^3 CT head data set is used and two sets of comparison are shown in Figure 8, where the upper set compares images generated with 512 samples and with 640 samples, and the lower set compares images generated with 1280 samples and 1408 samples. Figure 8 (c) shows that the differences between Figure 8 (a) and (b) are quite noticeable, and the percentage of red pixels is 39.54%. From Figure 8 (f), we can find that the differences between Figure 8 (d) and (e) are not so noticeable, where the percentage of red pixels is 28.46%.

4.5 Discussions

The experimental results show that differences between two direct volume rendered images are quite noticeable in transparent regions, indicating that different DVR algorithms or the same algorithm with different specifications are sensitive to these regions because of the inner structures of the data visualized there. Thus the choice of DVR algorithms or specifications have major impact on the visual result of the transparent regions. For shading methods, the visual appearance of images generated with pre-shaded and post-shaded are quite different as shown in the VDP result (Figure 5 (c)). The image quality of separate color interpolation is considered having color-bleeding artifacts [28]. The VDP result provides a distribution of such artifacts that may be noticed in the regions with red pixels. For gradient estimation schemes, the intermediate difference operator in DVR offers a better shading of the images [29] and the VDP result indicates such differences clearly. For sampling rates, the differences between two images with higher sampling rates are less than those two images with lower sampling rates. With enough high sampling rates, further increasing the sampling rate may not improve the image quality. With our framework, the differences between direct volume rendered images can be easily identified quantitatively. Thus it may be used for researchers to determine an appropriate DVR algorithm or a set of specifications for their research and applications.

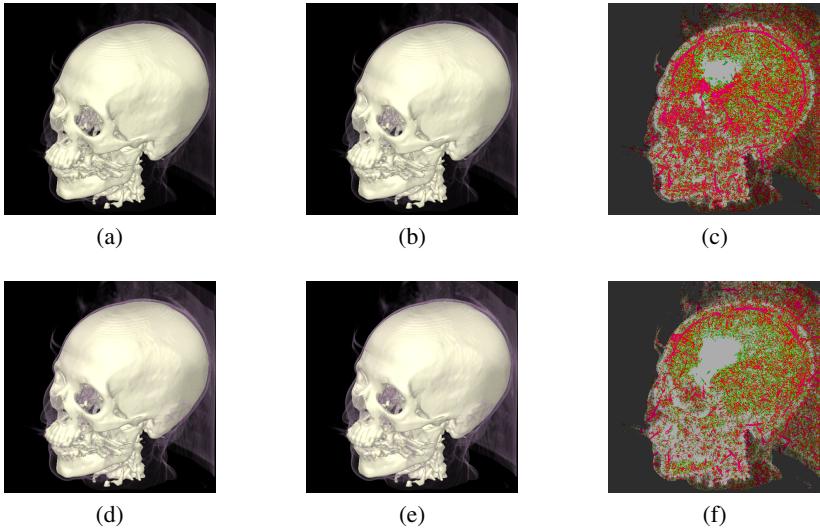


Fig. 8. Comparison of direct volume rendered images with different sampling rates (256^3 CT head): (a) Image rendered with 512 samples; (b) Image rendered with 640 samples; (c) VDP result of (a) and (b) (Red pixels: 39.54%); (d) Image rendered with 1280 samples; (e) Image rendered with 1408 samples; (f) VDP result of (d) and (e) (Red pixels: 28.46%)

5 Conclusions and Future Work

In this paper, we proposed a framework for comparing direct volume rendered images generated by different algorithms or the same algorithm with different specifications. Two most popular DVR algorithms, GPU-based ray-casting and 3D texture slicing, are selected for comparisons. Some specifications including shading methods (pre-shaded *v.s.* post-shaded, separate color interpolated *v.s.* opacity-weighted color interpolated), gradient estimation schemes (central and intermediate difference operators), and sampling rate are also compared. The experimental results with real data sets show that we can get quantitative and perceptual comparison results with our framework. To conclude, this study is our first attempt to apply perception knowledge on direct volume rendered images. It will allow scientists and engineers to better understand volume data.

In the future, we would like to perform a psychophysical validation of VDP for DVR applications and use our framework to conduct a more comprehensive study involving more direct volume rendered images generated by different kernels such as different filters and optical models. As the computation of VDP is quite expensive due to the multiscale spatial processing involved in some of its components, we plan to implement VDP on GPUs and integrate it with existing GPU-based volume rendering algorithms into our framework to provide fast feedbacks of comparison results. In addition, VDP may be used for level-of-detail (LOD) selection in large volume visualization as what Wang et al. [30] have done recently. This fast comparative framework can then be used for evaluating direct volume rendered images from a perceptual point of view and steering the DVR process.

Acknowledgments. This work is supported by the Science and Technology Development Fund of Macao SAR grant 001/2005/A.

References

1. Levoy, M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* **8** (1988) 29–37
2. Westover, L.: Footprint evaluation for volume rendering. In: *Computer Graphics (ACM SIGGRAPH '90)*. Volume 24. (1990) 367–376
3. Lacroute, P., Levoy, M.: Fast volume rendering using a shear-warp factorization of the viewing transformation. In: *Proceedings of ACM SIGGRAPH '94*. (1994) 451–458
4. Rezk-Salama, C., Engel, K., Bauer, M., Greiner, G., Ertl, T.: Interactive volume rendering on standard pc graphics hardware using multi-textures and multi-stage rasterization. In: *Proceedings of ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware '00*. (2000) 109–118
5. Gelder, A.V., Kim, K.: Direct volume rendering with shading via 3d textures. In: *Proceedings of ACM/IEEE Symposium on Volume Visualization '96*. (1996) 22–30
6. Krüger, J., Westermann, R.: Acceleration techniques for gpu-based volume rendering. In: *Proceedings of IEEE Visualization '03*. (2003) 287–292
7. Xue, D., Crawfis, R.: Efficient splatting using modern graphics hardware. *Journal of Graphics Tools* **8** (2003) 1–21
8. Kaufman, A., Mueller, K.: Overview of volume rendering. In: *The Visualization Handbook*, Edited by C. D. Hansen and C. R. Johnson, Elsevier Butterworth-Heinemann (2005) 127–174
9. Williams, P.L., Uselton, S.P.: Metrics and generation specifications for comparing volume-rendered images. *Journal of Visualization and Computer Animation* **10** (1999) 159–178
10. Kim, K., Wittenbrink, C.M., Pang, A.: Extended specifications and test data sets for data level comparisions of direct volume rendering algorithms. *IEEE Transactions on Visualization and Computer Graphics* **7** (2001) 299–317
11. Pommert, A., Höhne, K.H.: Evaluation of image quality in medical volume visualization: The state of the art. In: *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention '02*. (2002) 598–605
12. Daly, S.: The visible differences predictor: An algorithm for the assessment of image fidelity. In: *Digital Image and Human Vision*, Edited by A. B. Watson, The MIT Press (1993) 179–206
13. Gaddipatti, A., Machiraju, R., Yagel, R.: Steering image generation with wavelet based perceptual metric. *Graphics Graphics Forum (EUROGRAPHICS '97)* **16** (1997) 241–251
14. Machiraju, R., Yagel, R.: Reconstruction error characterization and control: A sampling theory approach. *IEEE Transactions on Visualization and Computer Graphics* **2** (1996) 364–378
15. Mei'sner, M., Huang, J., Bartz, D., Mueller, K., Crawfis, R.: A practical evaluation of popular volume rendering algorithms. In: *Proceedings of IEEE Symposium on Volume Visualization '00*. (2000) 81–90
16. Ramasubramanian, M., Pattanaik, S.N., Greenberg, D.P.: A perceptually based physical error metric for realistic image synthesis. In: *Proceedings of ACM SIGGRAPH '99*. (1999) 73–82
17. Bolin, M., Meyer, G.: A perceptually based adaptive sampling algorithm. In: *Proceedings of ACM SIGGRAPH '98*. (1998) 299–309
18. Myszkowski, K.: The visible differences predictor: Applications to global illumination problems. In: *Proceedings of EUROGRAPHICS Workshop on Rendering '98*. (1998) 223–236

19. Rushmeier, H., Ward, G., Piatko, C., Sanders, P., Rust, B.: Comparing real and synthetic images: Some ideas about metrics? In: Proceedings of EUROGRAPHICS Workshop on Rendering '95. (1995) 82–91
20. Lu, A., Morris, C.J., Taylor, J., Ebert, D.S., Hansen, C., Rheingans, P., Hartner, M.: Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics* **9** (2003) 127–138
21. Ebert, D.S.: Extending visualization to perceptualization: The importance of perception in effective communication of information. In: *The Visualization Handbook*, Edited by C. D. Hansen and C. R. Johnson, Elsevier Butterworth-Heinemann (2005) 771–780
22. Interrante, V.: Art and science in visualization. In: *The Visualization Handbook*, Edited by C. D. Hansen and C. R. Johnson, Elsevier Butterworth-Heinemann (2005) 781–806
23. Chalmers, A., Cater, K.: Exploiting human visual perception in visualization. In: *The Visualization Handbook*, Edited by C. D. Hansen and C. R. Johnson, Elsevier Butterworth-Heinemann (2005) 807–816
24. Zhou, H., Chen, M., Webster, M.F.: Comparative evaluation of visualization and experimental results using image comparison metrics. In: Proceedings of IEEE Visualization '02. (2002) 315–322
25. Lubin, J.: A visual discrimination model for imaging system design and evaluation. In: *Vision Models for Target Detection and Recognition*, Edited by E. Peli, World Scientific Publishing (1995) 245–283
26. Li, B., Meyer, G., Klassen, R.: A comparison of two image quality models. In: *Human Vision and Electronic Imaging III*, SPIE Vol. 3299 (1998) 98–109
27. Wilson, H.R.: Psychophysical models of spatial vision and hyperacuity. In: *Spatial Vision*, Vol. 10, *Vision and Visual Dysfunction*, Edited by D. Regan, MIT Press (1991) 64–86
28. Wittenbrink, C.M., Malzbender, T., Goss, M.E.: Opacity-weighted color interpolation for volume sampling. In: Proceedings of Symposium on Volume Visualization '98. (1998) 135–142
29. van Scheltinga, J.T., Bosma, M., Smit, J., Lobregt, S.: Image quality improvements in volume rendering. In: Proceedings of the Fourth International Conference on Visualization in Biomedical Computing (VBC '96). (1996) 87–92
30. Wang, C., Garcia, A., Shen, H.W.: Interactive level-of-detail selection using image-based quality metric for large volume visualization. *IEEE Transactions on Visualization and Computer Graphics* (accepted) (2006)