Ambient Volume Scattering — Supplemental Results

Marco Ament, *Student Member, IEEE*, Filip Sadlo, *Member, IEEE*,
and Daniel Weiskopf, *Member, IEEE Computer Society*

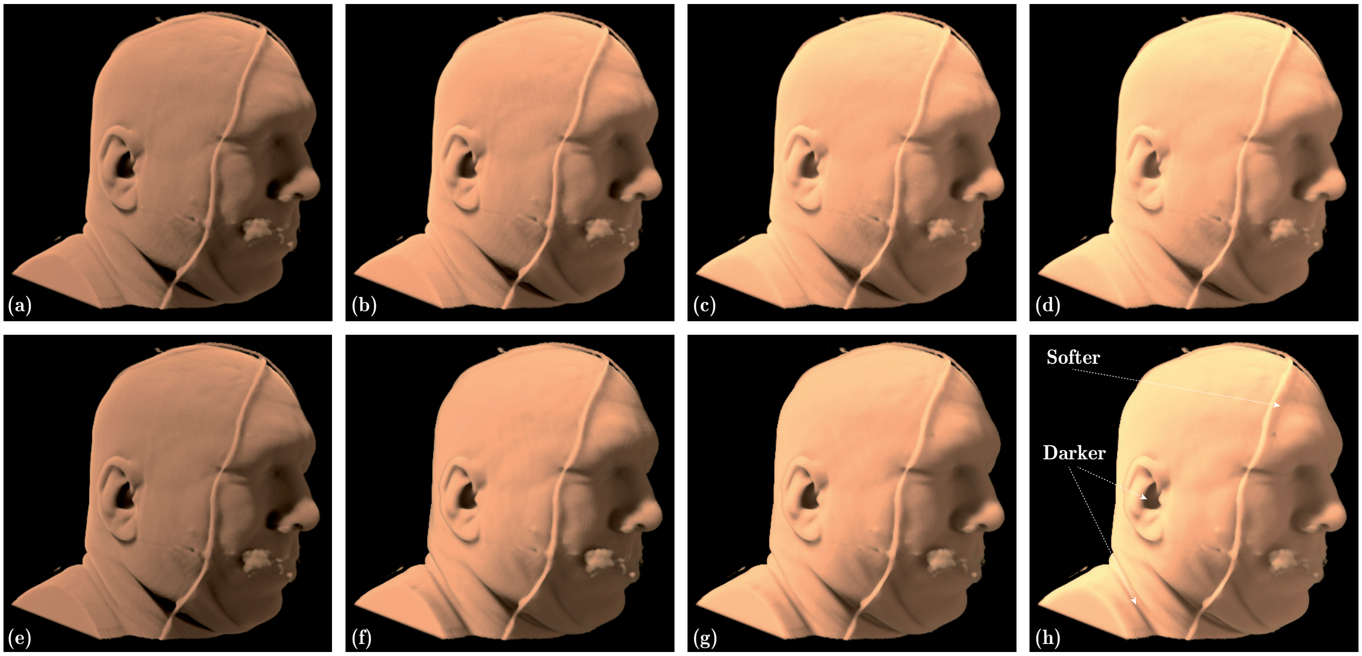

Fig. 1. Comparison of volumetric path tracing (top row) with ambient volume scattering (bottom row). Path tracing is terminated deterministically after (a) 1, (b) 2, and (c) 4 bounces. (d) Unbiased path tracing using Russian roulette for stochastic termination. (e) Single scattering with ray casting. Ambient volume scattering with an ambient radius of (f) 2, (g) 6, and (h) 8 voxels.

1 Comparison with Full Global Illumination

In this supplementary section, we qualitatively compare the illumination effects and the bias of ambient volume scattering with full global volume illumination. For comparison, we employ a CUDA-accelerated implementation of path tracing for participating media [1] with light source sampling to solve the radiative transfer equation accurately, using the same hardware as in the main paper (Intel Core i7 2.8 GHz CPU, 8 GB RAM, NVIDIA GTX-560 GPU). All ambient scattering results do neither include gradient-based shading nor ambient occlusion.

In Figure 1, we employ the Visible Human data set with an isotropic phase function ($g = 0.0$) and an albedo of $\Lambda = 0.9$ to compare subsurface scattering and shadow effects in an optically thick medium. Figures 1(a)–(d) are rendered with path tracing using 1 k samples per pixel and still contain some remaining noise. With a viewport size of 512^2 pixels, the render times are 6:36, 7:24, 8:30, and 12:02 minutes, respectively. The series of images of path tracing in Figures 1(a)–(d) shows the effect of increasing the number of scattering events (bounces) during light transport, all the way to unbiased global illumination in Figure 1(d). Intuitively, the ambient radius of ambient volume scattering controls the range of scattering effects in the neighborhood of each sample. Therefore, we study how varying radii can mimic the varying number of bounces that light undergoes in a participating medium. Figure 1(e) is rendered with ray casting and the single scattering model. The most important visual effects of the different scattering orders are reproduced by ambient scattering in Figures 1(f)–(h) by increasing the ambient radius. However, for larger radii, bias increases, since the assumption of a homogeneous neighborhood becomes more and more invalid. In Figure 1(h), the shadows on the forehead appear softer compared to path tracing due to overestimated scattering. In contrast, the skin fold close to the neck and the inner ear appear darker, since far-range scattering effects are not supported. By using the transmittance cache for single scattering and the radiance cache for ambient volume scattering, rendering speed is 110 fps for Figure 1(e) and 60 fps for Figures 1(f)–(h) with a viewport size of 512^2 pixels.

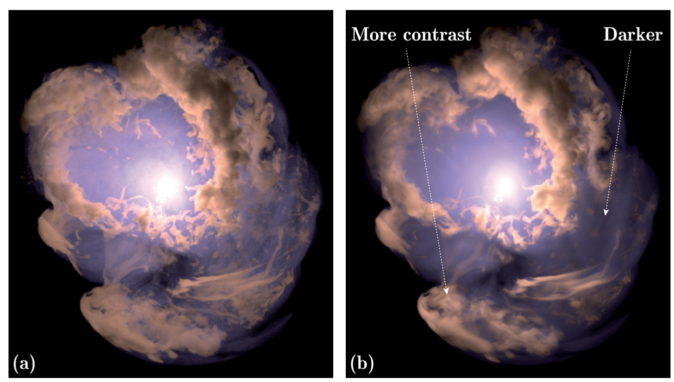

Fig. 2. (a) Unbiased path tracing using Russian roulette for termination. (b) Ambient volume scattering using a radius of 15 voxels.

In Figure 2, we employ the Supernova data set with an anisotropic phase function ($g = 0.6$) and an albedo of $\Lambda = 0.9$ to compare light transport in an optically thin and strongly heterogeneous medium. Figure 2(a) is rendered with unbiased path tracing using 1 k samples per pixel in 27:46 minutes. Figure 2(b) is rendered with ambient volume scattering and a radius of 15 voxels at 20 fps. We can observe two noticeable visual effects. First, the filaments in the lower nebula cluster of Figure 2(b) have a higher contrast compared to path tracing, since only light from ambient structures interacts with the medium. Second, some of the small features in the background appear darker because far-range scattering transport is not supported.

In visualization, the absence of global scattering transport is not as harmful as in photorealistic rendering. In fact, the perception of volumetric structures can be more difficult with full global illumination, because light arrives from all length scales and directions in the presence of a highly scattering medium. The strong interaction causes a visual diffusion of details, since even shadowed regions receive and reflect a significant amount of light. In the final image, the accumulated scattering events tend to reduce the contrast, which disturbs the perception of relative size and spatial depth. Ultimately, this is in conflict with the original goal of illumination in volume visualization. A similar observation can be found when comparing shadows from ambient occlusion with shadows from global illumination. Apart from the much higher performance, ambient occlusion often produces more distinct and pleasant shadows with more contrast than with an expensive Monte-Carlo simulation.

The mesoscopic model of ambient scattering prevents such exaggerated light transport naturally, but the missing contributions can also underestimate illumination, which results in darker features and can hide important information compared to global illumination. However, these structures can be enhanced easily by adding small contributions of self-emission or ambient light, e.g., by means of ambient occlusion. In this way, the visualization can be better adapted to the user's requirements and the rendering parameters are easier to control than with global illumination.

References

[1] H. E. Rushmeier. *Realistic image synthesis for scenes with radiatively participating media.* PhD thesis, Cornell University, 1988.