

A Zero-Variance-Based Sampling Scheme for Monte Carlo Subsurface Scattering

Jaroslav Krivánek*
Charles University in Prague

Eugene d'Eon
Weta Digital

We present a new variance reduction technique for unbiased Monte Carlo simulation of subsurface scattering—an important rendering problem for translucent materials like skin. The method is derived from the theory of zero-variance random walks, so far unexploited in computer graphics. A wealth of other potential applications of zero-variance estimators in rendering is yet to be explored.

We simulate subsurface scattering using a Monte Carlo random walk in the volumetric medium under the surface. The simulation involves two steps: transition distance sampling and direction sampling. Traditionally, the pdfs used for this purpose emulate the underlying physical processes (exponential law for distance sampling $p_d(s) = \sigma_t e^{-s\sigma_t}$, phase function $p_{ph}(\omega_o|\omega_i)$ for direction sampling). However, this sampling is purely local as it has no information about where the important parts of the entire domain are. In subsurface scattering simulation it is not useful to explore the medium far from the boundary, given that we are interested in paths that make it back out of the medium. This idea can be formalized using the notion of the importance function: A zero-variance estimator can be constructed by sampling paths proportionately to the product of the importance function and the classical pdfs [Hoogenboom 2008], and an approximation of this zero-variance ideal yields estimators with low variance. Dwivedi [1982] exploited this idea in deep-penetration transport problems such as reactor shielding. We show how this work can be adapted to reduce variance of subsurface scattering simulation.

Importance function. In Monte Carlo subsurface scattering simulation, the important event is emerging back from the medium, so the outside acts as a uniform source of importance. The importance function approximation within the medium is formed from the discrete eigenfunction of the exact transport solution for the homogeneous half-space problem [McCormick and Kuscer 1973]. It is a simple exponential of optical depth τ from the boundary, $I(\tau) \sim e^{-\tau/\nu_0}$, where ν_0 is the largest discrete eigenvalue of the transport operator, which, for isotropic scattering, is the positive real solution of $1 = \alpha\nu_0 \tanh^{-1} \frac{1}{\nu_0}$. Note that ν_0 only depends on the single-scattering albedo α . The normalized product of the importance function $I(\tau)$ and the classical pdfs yields the ‘guided’ sampling pdfs that prefer sampling paths toward the boundary. These pdfs are given below.

Direction sampling. Sampling of the outgoing direction after a collision is factorized into azimuthal angle selection (which is uniformly distributed in $[0, 2\pi)$) and the direction cosine sampling. The classical sampling in an isotropic medium selects a direction cosine u_o for the outgoing direction uniformly from $[-1, 1]$, so the pdf is $p(u_o) = 1/2$. Dwivedi’s procedure leads to the the following distribution for outgoing direction cosine u_o (measured from the principal axis perpendicular to the half-space medium boundary): $p'(u_o) = \frac{\alpha}{2} \frac{\nu_0}{\nu_0 - u_o}$.




Figure 1: Our method lowers variance per sample while increasing performance in Monte Carlo subsurface scattering simulation. The classical sampling results (left) use 100 samples/pixel. Our results use equal time with about 50% more samples/pixel (right). The two rows show results for two different env. maps.

Distance sampling. Given an outgoing direction cosine u_o the transition distance s is chosen from an exponential distribution with a modified extinction coefficient $\sigma'_t = \sigma_t(1 - u_o/\nu_0)$. The distance sampling pdf $p'_d(s) = \sigma'_t e^{-s\sigma'_t}$ can be sampled as $s = -\log(\xi_s)/\sigma'_t$ where ξ_s is a uniform random variate.

Because the eigenvalue $\nu_0 \rightarrow \infty$ as $\alpha \rightarrow 1$, we see that the Dwivedi pdfs approach the classical sampling for no absorption. The higher the absorption (the lower α) the more upward directions are preferred ($u_o > 0$) and the more transition distances are shortened when going deep and increased when heading towards the surface.

Application and results. Our implementation uses the surface normal at the point where a light path enters the medium as the principal axis perpendicular to the boundary (the direction cosine u_o is measured from this axis). To maintain the robustness of the method for highly curved surfaces, we use multiple importance sampling to combine Dwivedi’s and classical sampling pdfs. The result is a simple unbiased Monte Carlo sampling method for subsurface scattering that avoids the approximations incurred by the previously used BSSRDFs, and that is easy to incorporate in existing MC rendering algorithms. A result generated using our method implemented in a spectral path tracer is shown in Figure 1 (additional results are in the supplemental document). Our method significantly reduces variance per sample and—as a bonus—generates paths that are shorter on average, thereby increasing performance. We believe that the zero-variance theory, which is at the basis of the present method, opens up numerous exciting opportunities for future work.

Acknowledgments. Part of this work was done while the first author was with Weta Digital. The work was partially supported by the Czech Science Foundation grant P202-13-26189S.

References

- CHRISTENSEN, P. 2003. Adjoints and importance in rendering: an overview. *IEE TVCG* 9, 3 (July), 329–340.
- DWIVEDI, S. 1982. A new importance biasing scheme for deep-penetration Monte Carlo. *Annals of Nuclear Energy* 9, 7.
- HOOGENBOOM, E. 2008. Zero-variance monte carlo schemes revisited. *Nuclear science and engineering* 160, 1, 1–22.
- MCCORMICK, N., AND KUSCER, I. 1973. Singular eigenfunction expansions in neutron transport theory. *Advan. Nucl. Sci. Technol.*, v. 7, pp. 181–282 7.

*e-mail: jaroslav.krivanek@mff.cuni.cz