(Talk 0263) Smoother Subsurface Scattering

Ivan Neulander
Rhythm & Hues Studios

The irradiance-cached dipole diffusion approximation proposed by [Jensen and Buhler 2002] is a practical tool for capturing the subsurface scattering of many common materials in the context of film-quality rendering. One drawback of this model is the high irradiance point density it requires to attenuate sampling noise at low mean free paths, which can be expensive in terms of time and memory. A clever smoothing technique by [Langlands and Mertens 2007] filters out this sampling noise by blending the noisy irradiance-cache-derived diffusion approximation with a noise-free precomputed diffusion integral that is weighted by local irradiance. But this smoothing leaves an artifact at low sample densities, as shown in Figure 1: abrupt changes in irradiance, e.g. across shadow edges, tend to produce hard edges in the rendered image, irrespective of the mean free path. Moreover, increasing the smoothing radius hardens rather than softens these edges.

We extend the work of [Langlands and Mertens 2007] by adding two enhancements: 1) higher fidelity of shadow edges, at the expense of adding a small amount of noise; 2) the ability to interpolate among a sparse set of precomputed diffusion integrals, allowing continuous variation of the mean free path.

Figure 1: Results using 23k irradiance points: (a) no smoothing; (b) Langlands & Mertens smoothing; (c) our smoothing; (d) ground truth. See supplemental notes for timing data on this and other images.

Shadow Edge Fidelity

The [Langlands and Mertens 2007] technique separates the diffusion component of the BSSRDF integral into a **local term** $\int w(r)R_d(r)E(\mathbf{x})\,d\mathbf{x}$ and a **global term** $\int[1-w(r)]R_d(r)E(\mathbf{x})\,d\mathbf{x}$, where $w(r)$ is a weighting kernel defined so that $w(0)=1$ and $w(\alpha\cdots\infty)=0$ for a chosen smoothing radius of α. The global term is computed only from cached irradiance samples, as prescribed by [Jensen and Buhler 2002], while the local term is accurately precomputed using a constant irradiance value and multiplied by the actual irradiance at shading point $\mathbf{x}$. This arrangement supplants the noisy global term with the noise-free local term in the vicinity of irradiance-cached sample points. But because the local term varies linearly with local irradiance, rapid changes in the latter (within a radius of α) spill over into the overall radiant exitance. As the irradiance cache density increases, α naturally diminishes and the problem abates. However, as shown in Figure 2, the density required for passable shadow edges may be so high as to obviate the need for smoothing.

Our solution is to modify the irradiance value that drives the local term: in addition to local irradiance, we include a contribution from the entire irradiance cache. Conceptually, we insert the local irradiance point into the cache, using the appropriate weighing factor of $R_d(0)$ and an area weight equal to the average irradiance-cached sample area. Then we query this modified cache for the irradiance to apply to the local term. This query is identical to that for the global term, except for the extra point at $r=0$ and the differing kernel weights. Our implementation performs both queries in a single traversal of the octree, with essentially no redundant computation.

By incorporating cached irradiance into the local term, we reintroduce some variability into the radiant exitance. However, because R_d peaks sharply at zero, the resulting noise is inconspicuous except where the local irradiance is low. In effect, we replace hard shadow edges with noisier but appropriately soft ones.

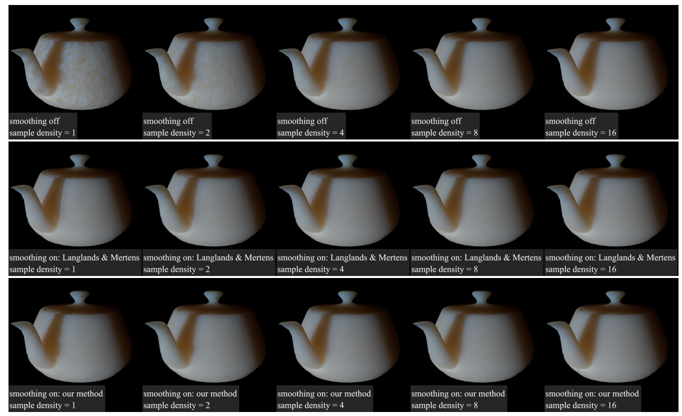

Figure 2: In the middle three columns, our method yields the best image. In the left column all three methods work poorly, while in the right all three are close in quality.

Continuous Mean Free Path Variation

The precomputed integral for the local term in [Langlands and Mertens 2007] varies nontrivially with the mean free path ℓ_u, so rendering multiple ℓ_u values entails multiple integrations. This precludes efficient rendering when ℓ_u varies continuously.

We lift this limitation by evaluating and caching the integral on demand at exponentially distributed values of ℓ_u, and linearly interpolating between adjacent integrals. As ℓ_u approaches zero, we limit the evaluation rate based on a user-defined threshold. We have found empirically that caching at integer powers of 2 is adequate, as illustrated in Figure 3. In general, our interpolated smoothing incurs a runtime overhead of under 5% relative to no smoothing.

Figure 3: Texture-mapped ℓ_u variation with our smoothing.

References

JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. In *SIGGRAPH '02 Proceedings*, ACM, 576–581.

LANGLANDS, A., AND MERTENS, T. 2007. Noise-free bssrdf rendering on the cheap. In *SIGGRAPH '07 posters*, ACM, 182.