Practical and Controllable Subsurface Scattering for Production Path Tracing

Matt Jen-Yuan Chiang^{*} Peter Kutz Brent Burley Walt Disney Animation Studios

Keywords: subsurface scattering, path tracing

Concepts: •Computing methodologies \rightarrow Rendering; Reflectance modeling;

1 Background

Subsurface scattering is ubiquitous in digital worlds created for films. While diffusion-based approximations are still widely used, they can produce undesirable and implausible results when the "semi-infinite slab" assumption is violated. The brute force alternative, Monte Carlo simulation of volumetric path tracing, is much more robust over geometric variations. However, its physically based nature makes it unintuitive for artists. In this talk, we exploit the best of both worlds: taking advantage of the robustness of the path-traced approach while allowing artists to use the production-proven intuitive controls from the Normalized Diffusion model [Burley 2015]. We also introduce a sampling scheme that improves the performance of the brute force approach to be practical for production rendering.

2 Intuitive Parameterization

With Normalized Diffusion, artists specify the apparent color of the surface, A, along with a scattering distance, d. These parameters directly control the diffusion profile in an intuitive way. In contrast, volumetric path tracing requires single-scattering parameters, and the apparent color, i.e. the multiple-scattering albedo, is the end result of numerous scattering events. Unfortunately, it is impossible to analytically invert the random walk process to infer the parameters describing individual scattering events from the multiple-scattering result.

To simplify things, as with diffusion, we assume isotropic scattering with a diffuse interface. Our insight is that, on a semi-infinite slab, the single-scattering albedo, α , is then the only factor affecting the multiple-scattering albedo A while the extinction coefficient, σ_t , merely scales the path lengths and hence only affects the translucency. To map the relationship between A and α , we rendered a slab with differing single scattering albedos and fixed extinction coefficient, illuminated by a white environment. The resulting data resembles an exponential relationship which we fit using a third degree polynomial in log-space. We infer the extinction coefficient from our scattering distance using the mapping from Christensen et al. [2015]. Combining these two mappings gives us our reparameterization:

$$\alpha = 1 - e^{-5.09406A + 2.61188A^2 - 4.31805A^3} \tag{1}$$

$$s = 1.9 - A + 3.5 * (A - 0.8)^2$$
⁽²⁾

$$\sigma_t = 1/(d*s) \tag{3}$$

*e-mail:matt.chiang@disneyanimation.com

SIGGRAPH '16, July 24-28, 2016, Anaheim, CA,

DOI: http://dx.doi.org/10.1145/2897839.2927433



Figure 1: Translucent materials rendered using Normalized Diffusion (left) vs. path-traced subsurface scattering (right). With our parameterization, path-traced subsurface scattering shares the same controls of the production-proven Normalized Diffusion and produces identical results on a semi-infinite slab (top row). Yet it brings out more geometric details especially around highly curved areas.

3 Sampling

Sampling volumetric path tracing consists of two steps: directional and distance sampling. Directional sampling is trivial since we assume isotropic scattering. However, sampling distance is a challenge when the volume extinction is chromatic, which is usually the case. Taking skin for example, red light travels much further than other wavelengths before a statistical scattering/absorption event. The target integrand f(s) for sampling distance can be perfectly importance sampled according to the probability density $p(s) = \sigma_t e^{-\sigma_t s}$ for each wavelength. To reduce sample variance, we combine estimators for different wavelengths using MIS where each estimator is given a relative weight proportional to $f(s)/p(s) = \sigma_s/\sigma_t$, which is the single-scattering albedo. We also multiply the accumulated per-wavelength throughput of the current path into the weight reasoning that as the throughput goes to zero for some wavelength, no further samples are needed.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). © 2016 Copyright held by the owner/author(s).

ISBN: 978-1-4503-4282-7/16/07

4 Discussion

With our sampling scheme, brute force path tracing achieves rendering time and noise level comparable to ray-traced diffusion, at least with parameters typical of skin. While more rays generally need to be traced with path tracing, determining sample points for diffusion still requires a significant number of probe rays [King et al. 2013], and unlike the probe rays, the path-tracing segments are of fixed length making them generally less expensive to trace. We have integrated path-traced subsurface scattering with our BSDF [Burley 2015] by replacing the non-directional (i.e. Lambertian) diffuse term and applying the view- and light-dependent diffuse terms to the entry and exit points respectively. This new shading model has been successfully adopted by production at the Walt Disney Animation Studios.

5 Acknowledgements

Dragon model courtesy of the Stanford Computer Graphics Laboratory. All other images ⓒ Disney Enterprises, Inc.

References

- BURLEY, B., 2015. Extending Disney's physically based BRDF with integrated subsurface scattering. ACM SIGGRAPH 2015 Physically Based Shading Course Notes, July.
- CHRISTENSEN, P. H. 2015. An approximate reflectance profile for efficient subsurface scattering. In *ACM SIGGRAPH 2015 Talks*, ACM, 25.
- KING, A., KULLA, C., CONTY, A., AND FAJARDO, M. 2013. BSSRDF importance sampling. In *ACM SIGGRAPH 2013 Talks*, ACM, 48.