

Approximate Reflectance Profiles for Efficient Subsurface Scattering

Per H. Christensen

Brent Burley

Pixar Animation Studios

Walt Disney Animation Studios

Pixar Technical Memo #15-04 — July, 2015



Figure 1: Images rendered with the approximate subsurface scattering reflectance profiles presented in this report. © Disney/Pixar. (Prometheus statue modeled by Scott Eaton; head data courtesy of Infinite Realities via Creative Commons; alien, fruits, and candle rendered by Dylan Sisson; sheep modeled by Chris Scoville.)

Abstract

We present three useful parameterizations of a BSSRDF model based on empirical reflectance profiles. The model is very simple, but with the appropriate parameterization it matches brute-force Monte Carlo references better than state-of-the-art physically-based models (quantized diffusion and photon beam diffusion) for many common materials. Each reflectance profile is a sum of two exponentials where the height and width of the exponentials depend on the surface albedo and mean free path length. Our parameterizations allow direct comparison with physically-based diffusion models using the same parameters. The parameterizations are determined for perpendicular illumination, for diffuse surface transmission (where the illumination direction is irrelevant), and for an alternative measure of scattering distance. Our approximations are useful for rendering ray-traced and point-based subsurface scattering.

Keywords: Rendering, translucent materials, subsurface scattering, functional approximation, BSSRDF, reflectance profile.

1 Introduction

Realistic modeling of subsurface scattering is important for rendering believable images of translucent materials such as skin, meat,

fruits, plants, wax, marble, jade, milk and juice. Computer graphics researchers have developed increasingly sophisticated and accurate physically-based subsurface scattering models from the simple dipole diffusion model [Jensen et al. 2001] to the quantized diffusion model [d’Eon and Irving 2011] and photon beam diffusion and diffuse single-scattering model [Habel et al. 2013]. Here we introduce three parameterizations of an empirical model that is as simple as the dipole but matches brute-force Monte Carlo references better than even photon beam diffusion.

The main reasons for replacing physically-based models with an approximation are:

- No need for numerical inversion [Jensen and Buhler 2002; Habel et al. 2013] of the user-friendly surface albedo and mean free path length input parameters to less intuitive volume scattering and absorption coefficients.
- Built-in single-scatter term.
- Faster evaluation, simpler code and no need for lookup-tables.
- No ad-hoc correction factor $\kappa(r)$ to make the theory fit Monte Carlo references [Donner and Jensen 2007; Habel et al. 2013].
- A simple cdf for importance sampling.

2 Background and related work

2.1 Monte Carlo simulation and BSSRDFs

The most general method to compute subsurface scattering is to treat the object as a volume and run a brute-force Monte Carlo simulation [Kalos and Whitlock 1986; Wang et al. 1995]. However, this can be extremely slow, particularly for complex scenes.

The function that describes how light enters an object, scatters around inside it, and then leaves the object is the BSSRDF — the bidirectional surface scattering reflectance distribution function. Donner et al. [2009] used Monte Carlo particle tracing to tabulate an empirical BSSRDF model for a flat surface on a homogeneous semi-infinite volume. They represented the hemispherical distribution of light leaving the surface, depending on the angle of the incident light, the relative position of the exitant light, and the physical parameters (volume albedo, mean free path length, phase function, and index of refraction). Their tables took months to compute and contain around 250MB of data.

2.2 Physically based reflectance profiles

The BSSRDF S is often simplified as a product of a radially symmetric (1D) diffuse *reflectance profile* R , two directional Fresnel transmission terms F_t , and a constant C [Jensen et al. 2001; d’Eon and Irving 2011; Jimenez et al. 2015]:

$$S(x_i, w_i; x_o, w_o) = C F_t(x_i, w_i) R(|x_o - x_i|) F_t(x_o, w_o) \quad (1)$$

Figure 2 shows examples of reflectance profiles for various surface albedos; these curves were computed with Monte Carlo simulation. The vertical axis shows $rR(r)$ rather than $R(r)$ since $R(r)$ is always integrated radially over the surface (and has a sharp peak close to $r = 0$).




Figure 2: Reflectance profiles $R(r)$ for various surface albedos.

The dipole diffusion model [Jensen et al. 2001] is an approximation of subsurface scattering that has been diffused after many scattering events. This model is simple, fast to evaluate and widely used; however, it is also overly blurry and results in a waxy look. The scattering was parameterized by the volume scattering and absorption coefficients σ_s and σ_a (or, equivalently, the volume scattering albedo $\alpha = \sigma_s/\sigma_t = \sigma_s/(\sigma_s + \sigma_a)$) and volume mean free path length $\ell = 1/\sigma_t = 1/(\sigma_s + \sigma_a)$.

In follow-up work, Jensen and Buhler [2002] introduced a more intuitive parameterization of subsurface scattering: surface albedo (diffuse surface reflectance) $A = \int_0^\infty R(r) 2\pi r dr$ and diffuse

mean free path length ℓ_d on the surface. We use similar intuitive parameters in our approximations. (We use the symbol A for surface albedo instead of the commonly used R_d since R_d is also frequently used for the diffusion (multi-scattering) part of the reflectance profile.) d’Eon later presented a more physically accurate dipole diffusion model simply called “a better dipole” [d’Eon 2012]. Recently Frisvad et al. [2014] introduced a directional dipole model that lifts the assumption of a radially symmetric diffusion profile; this improves the accuracy for e.g. non-perpendicular (oblique) illumination of objects with smooth refractive surfaces (such as milk, fruit juice, and other liquids with suspended particles).

d’Eon and Irving [2011] introduced quantized diffusion: improved diffusion theory and an extended source term instead of just a dipole. This results in a more realistic look with sharper features, but is also much more complicated and time-consuming to calculate. The resulting diffuse reflectance profile was approximated as a sum of Gaussians. Our approximations are more accurate and vastly simpler to compute.

The photon beam diffusion paper by Habel et al. [2013] had three main contributions: a photon beam diffusion model that is as accurate as quantized diffusion but much faster to evaluate, an accurate diffuse single-scattering model, and elegant handling of oblique refraction into the material. Our approximations are faster and also more accurate than photon beam diffusion for symmetric scattering (either perpendicular illumination or ideal diffuse surface transmission). The photon beam model has an empirical term $\kappa(r)$ to make the theoretical results better match Monte Carlo references in the mid-distance range; our model abandons theory entirely and consists of purely empirical terms.

All the diffusion models above require separate handling of single scattering, either by explicit ray tracing or separate integration. Both methods are slow. Our model includes single scattering, so no expensive separate calculation is needed.

2.3 Approximate reflectance profiles

Reflectance profiles like the ones in Figure 2 can be represented with tables. For example, we could store tables of the reflectance profile as function of distance r for surface albedos $0, 0.01, 0.02, \dots, 1$, and then interpolate between them for a given surface albedo and distance.

However, our inspiration comes from the frequently used technique of approximating complex functions by simpler ones. A good example is the Fresnel reflection and refraction formulas. Fresnel reflection and refraction can be modeled based on physics (Maxwell’s equations and energy constraints) as a sum of two terms, one for perpendicular and one for parallel polarized light. But Schlick [1994] made the observation that the resulting curve can be closely approximated by a simple polynomial, and this approximation is widely used in computer graphics since it is simpler, faster to evaluate, and gives no visible difference. We would like a similar approximation for subsurface scattering.

The reflectance profile has been approximated reasonably well with a sum of zero-mean Gaussians [d’Eon et al. 2007; Yan et al. 2012; Jimenez et al. 2015], and approximated more crudely with a single Gaussian or a cubic polynomial [Kling et al. 2013].

Burley [2013; 2015] noted that the shape of the diffuse reflectance profile can be approximated quite well with a curve in the shape of a sum of two exponential functions divided by distance r :

$$R(r) = \frac{e^{-r/d} + e^{-r/(3d)}}{8\pi dr} \quad (2)$$

The d parameter shapes the height and width of the curve and can be set based on artistic preference or determined based on physical parameters. With this expression for $R(r)$, any positive value of d gives a surface albedo of 1, hence Burley named it *normalized diffusion*. By multiplying by surface albedo A and picking an appropriate value for d we can obtain a remarkably accurate fit for many common materials. This model is implemented in Walt Disney Animation Studio’s Hyperion renderer; Figure 3 shows an example of the use of this scattering model in the Disney movie “Big Hero 6”.



Figure 3: Image from “Big Hero 6” rendered with Hyperion.

In the following sections we present simple analyses of how to best scale and stretch the normalized diffusion curve to match Monte Carlo references over the full range of valid surface albedos. In other words, we determine suitable “translations” from physical parameters to d . This enables use of the same physical parameter for scattering distance as used by physically-based diffusion models, and facilitates direct comparison with those models.

3 Searchlight configuration

We first consider the so-called searchlight configuration where a focused beam of light is incident on a semi-infinite homogeneous medium under a flat surface — see Figure 4. Photons are transmitted through the surface, get scattered by the medium, and ultimately get absorbed or escape the medium back at the surface. The distribution of photons exiting the surface forms a reflectance profile $R(r)$, which is radially symmetric for normally-incident light or diffuse transmission.




Figure 4: Searchlight configuration.

In this section we assume that the photons first travel straight down perpendicular to the surface, a simplifying assumption also used in e.g. the MCML simulation package [Wang et al. 1995] and elsewhere. We also assume that the phase function is isotropic — anisotropic phase functions are often dealt with using similarity of moments by using a reduced scattering coefficient $\sigma'_s = (1 - g)\sigma_s$, but more on this under future work in section 8.

3.1 Monte Carlo references

Figure 5 shows reflectance profiles $R(r)$ for surface albedos between 0.1 and 0.9 with mean free path $\ell = 1$ and anisotropy $g = 0$. (These are actually the same data as in Figure 2 but now with a log vertical axis.) These reference profiles are computed with a brute-force Monte Carlo simulation similar to MCML [Wang et al. 1995] and are our target curves for approximation. The surface albedo is computed without the Fresnel terms as $A = \int_0^\infty R(r) 2\pi r dr$ with volume scattering and absorption coefficients σ_s and σ_a chosen such that the mean free path length $\ell = 1/(\sigma_s + \sigma_a)$ is 1 (i.e. $\alpha = \sigma_s, \sigma_a = 1 - \sigma_s = 1 - \alpha$) and the desired surface albedo A is reached when integrating $R(r)$.

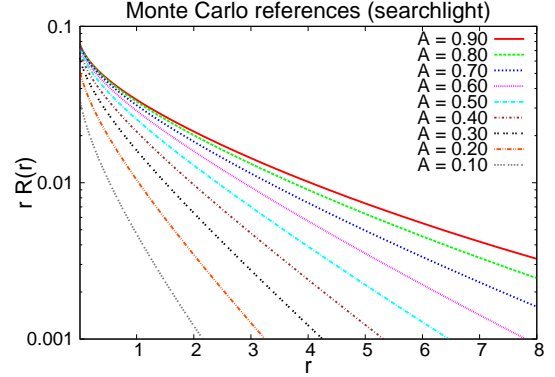


Figure 5: Reflectance profiles $R(r)$ for the searchlight configuration for surface albedos A between 0.1 and 0.9. (Log vertical axis, $\ell = 1, g = 0$.)

3.2 Functional approximation

We wish to determine the appropriate value of d corresponding to a physically meaningful quantity, and have chosen volume mean free path length ℓ . If we express the relationship between d and ℓ with a scaling factor s we can set $d = \ell/s$ in Equation 2 and get:

$$R(r) = A s \frac{e^{-sr/\ell} + e^{-sr/(3\ell)}}{8\pi\ell r} . \quad (3)$$

Next we note that to determine s by curve fitting it is sufficient to consider $\ell = 1$ since the *shape* of the Monte Carlo reference curve for a given A is independent of ℓ : $R(r, \ell) = R_{\ell=1}(r/\ell) / \ell^2$. So for curve fitting we only need to consider:

$$R_{\ell=1}(r) = A s \frac{e^{-sr} + e^{-sr/3}}{8\pi r} . \quad (4)$$

Now we need to determine good s values for the valid range of A . We use brute-force random sampling of the s parameter space to minimize relative error over r_i : $\sum_i \frac{|R(r_i) - R_{MC}(r_i)|}{R_{MC}(r_i)}$ to determine the optimal value s for any given value of A .

Figure 6 shows the fit of our reflectance profile approximation compared to Monte Carlo references, best fitting two Gaussians, photon beam diffusion plus single scattering, dipole diffusion plus single scattering, and better dipole diffusion plus single scattering for surface albedos 0.2, 0.5, and 0.8. The figures illustrate that two Gaussians cannot match the reference curves as well as our parameterization of normalized diffusion. It also shows that our approximation is closer to the MC reference points than the dipole, better dipole, and photon beam diffusion with single scattering.




Figure 6: Fit of various reflectance profile models for surface albedos $A = 0.2, 0.5,$ and 0.8 . (Corresponds to volume albedo α of $0.686, 0.938,$ and 0.9939 respectively, with $\ell = 1$. Log vertical axes.)

We can simply generate a table of s values for values of $A = 0.01, 0.02, \dots, 0.99$; some of these values are plotted as data points in Figure 7. If we use these values and interpolate for in-between surface albedos, we get 4.9% average relative error with respect to the Monte Carlo references. But here we present a simple function that is even more compact and easy to evaluate.




Figure 7: Data points and fitted curve for s .

With a bit of manual curve fitting we have found that the following simple expression for the scaling factor s gives a good fit to the optimal values:

$$s = 1.85 - A + 7|A - 0.8|^3. \quad (5)$$

This function is plotted as the curve in Figure 7. The relative error of the reflectance profile with respect to the Monte Carlo references is on average 5.5% over the full range of surface albedos with this expression for s . Compared to all the approximations and assumptions implicitly built into the Monte Carlo references (semi-infinite homogeneous volume, flat surface, searchlight configuration, etc.) this is actually a rather small error.

4 Diffuse surface transmission

In the previous section we assumed that the light enters straight into the volume in a direction perpendicular to the surface. In this section we will model the subsurface scattering reflectance profile after *ideal diffuse transmission* at the surface. This may be a more appropriate model for rough surface materials such as dry (non-sweaty) skin, make-up, most fruits, and rough (unpolished) marble, and also for the situation where we ignore (or do not know) the direction of the incoming light.

Figure 8 shows reflectance profiles $R(r)$ for subsurface scattering after ideal diffuse surface transmission (cosine distribution), again computed with Monte Carlo simulation. The general shape of the

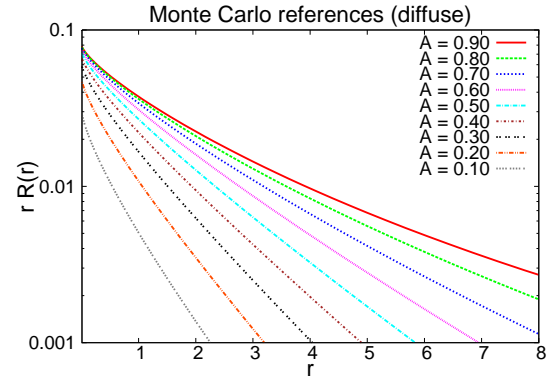


Figure 8: Reflectance profiles $R(r)$ for diffuse transmission for surface albedos A between 0.1 and 0.9 . (Log vertical axis, $\ell = 1, g = 0$.)

reflectance profiles for this case is similar to the searchlight case, so we use the same functional approximation, Equation 3, but compute new s values. For optimal s values we get only 2.6% average relative error with respect to the Monte Carlo references.

Again using manual curve fitting we have found that this expression for s gives a good fit to the optimal values:

$$s = 1.9 - A + 3.5(A - 0.8)^2. \quad (6)$$

Figure 9 shows data points and the fitted curve for s . Using this expression for s , the average relative error with respect to the Monte Carlo references is only 3.9%.




Figure 9: Data points and fitted curve for s .

In practical use, we have found that there is not much difference between the look of the searchlight approximation and the diffuse-

transmission approximation. So even though these abstractions are intended to represent two very different classes of surfaces, in practical VFX and CG animation work this distinction might not be particularly important.

5 dmfp as parameter

We now return to the searchlight configuration. It is possible to use an alternative parameterization of the scattering distance: diffuse mean free path (dmfp) on the surface, ℓ_d , instead of the mean free path (mfp) in the volume, ℓ . To calculate the diffuse mean free path length corresponding to σ_s and σ_a we can first compute the diffusion coefficient

$$D = (\sigma_t + \sigma_a) / (3\sigma_t^2) . \quad (7)$$

Given D we can then compute the effective transport extinction coefficient $\sigma_{tr} = \sqrt{\sigma_a/D}$ and then $\ell_d = 1/\sigma_{tr}$. In order to compute Monte Carlo reference curves we simply determine which pair of σ_s and σ_a values give the desired A value and also give $\ell_d = 1$: choose the ratio of σ_s, σ_a that gives A and then scale both σ_s and σ_a together to get $\ell_d = 1$. We have found that a good fit to these curves can be obtained by replacing ℓ with ℓ_d in Equation 3 and using this simple expression for s :

$$s = 3.5 + 100(A - 0.33)^4 . \quad (8)$$

Figure 10 shows data points and the fitted curve for s . The average error is 6.4% with optimal s values and 7.7% using the expression above.




Figure 10: Data points and fitted curve for s .

6 Practical detail: importance sampling

For importance sampling proportional to the radial distance between entry point and exit point, we need the cdf (cumulative distribution function) corresponding to the product of $R(r)$ and $2\pi r$. For physically-based BSSRDFs this cdf has to be computed with numerical integration, which is cumbersome and degrades performance. But fortunately Burley’s normalized diffusion $R(r)$ times $2\pi r$ is easily integrated, the cdf is:

$$\text{cdf}(r) = \frac{\int_0^r R(t) 2\pi t dt}{\int_0^\infty R(t) 2\pi t dt} \quad (9)$$

$$= \frac{A/4 (4 - e^{-r/d} - 3e^{-r/(3d)})}{A} \quad (10)$$

$$= 1 - \frac{1}{4}e^{-r/d} - \frac{3}{4}e^{-r/(3d)} \quad (11)$$

This simple cdf can be used with both our parameterizations: $d = \ell/s$ or $d = \ell_d/s$.




Figure 11: (a) Skin-colored materials rendered with photon beam diffusion and single-scattering. (b) Same materials rendered with our approximate reflectance profile.

The inverse cdf is $r = \text{cdf}^{-1}(\xi)$ where ξ is a random variable between 0 and 1. Unfortunately the cdf is not analytically invertible, but we can handle this in (at least) three ways: 1) randomly pick one of the two exponents, use its inverse as a cdf, and then weigh the results using MIS (multiple importance sampling); 2) a few Newton iterations; 3) scale a precomputed table of cdf^{-1} for $d = 1$ by d .

7 Discussion and results

Using these simple formulas is many times faster than computing quantized diffusion or photon beam diffusion with single-scattering — all requiring numerical integration. In fact, our formulas are even faster than the simple dipole diffusion model. Since the parameters are the same, switching from one of the physically-based diffusion models to our approximate diffusion model is simple.

In practical use of the physically-based models, they can be tabulated on-demand during rendering and the tables interpolated during the following lookups, so in practice the run-time saving using our approximation is actually rather small. Nonetheless, the code simplification is significant from a rendering author point-of-view. Particularly the code to numerically invert surface albedo and mean free path length to volume scattering and absorption coefficients is not entirely trivial. Eliminating that code plus the optimization code for table generation and lookups, and replacing it all with our very simple formulas for s and normalized diffusion $R(r)$ is a useful and practical simplification. As a bonus, the cdf for importance sampling is very simple.

Our approximate diffusion model is useful for ray-traced and point-based subsurface scattering, and the searchlight mfp and dmfp parameterizations are implemented as two of the subsurface scattering models in Pixar’s RenderMan renderer.

Figure 11 shows a surface illuminated by a spot light; the surface consists of three strips with different skin-like surface albedos. Photon beam diffusion plus single-scattering on the left and a BSSRDF using our approximate reflectance profile (searchlight configuration from section 3) with the same parameters on the right.

Figure 12 shows a close-up of a human head with photon beam diffusion plus single-scattering (with scattering straight into the surface) on the left and with a BSSRDF using our approximate reflectance profile with the same parameters on the right. Notice the characteristic reddish glow of subsurface scattering through the back-lit earlobe in both images.

Figure 1 shows other examples of subsurface scattering rendered with our approximate reflectance profiles: marble, human skin, plastic, fruits, candle wax, and animal skin.




Figure 12: Head close-up. (a) Rendered with photon beam diffusion and single-scattering. (b) Rendered with our approximate reflectance profile using the same scattering parameters — surface albedo texture and mean free path lengths. (Head data courtesy of Infinite Realities via Creative Commons.)

8 Conclusion and future work

We have presented three parameterizations of an approximate reflectance profile for simple and efficient rendering of subsurface scattering. The relative error compared to Monte Carlo references is on average 5.5%, 3.9%, and 7.7% respectively over the full range of surface albedos. We firmly believe that aiming for higher accuracy would be wasted while still retaining the much more serious infinite slab assumption (semi-infinite homogeneous volume with flat surface).

Future work includes looking at a less restrictive set of assumptions. For example, it would be very interesting to try to fit simple approximations to the diffusion profiles for oblique incident directions and non-radially-symmetric scattering [Donner et al. 2009; Habel et al. 2013; Frisvad et al. 2014; d’Eon 2014]. It might be as simple as having the scale s depend on the polar and relative azimuthal angle of the incident illumination (as well as surface albedo). Also, diffuse subsurface scattering materials with anisotropic scattering ($g \neq 0$) are often approximated using similarity of moments: $\sigma'_s = (1 - g)\sigma_s$. This may be a reasonable assumption after many scattering events, but for single-scattering and low-order scattering it is probably not. It would be interesting to fit curves to Monte Carlo references computed with anisotropic scattering.

Acknowledgements

Many thanks to our colleagues in the RenderMan and Hyperion teams for their support. Also thanks to Christophe Hery and Ryusuke Villemin at Pixar and Wojciech Jarosz and Ralf Habel at Disney for many productive discussions about subsurface scattering, and to Dylan Sisson for images.

References

BURLEY, B. 2013. Subsurface scattering investigation. Unpublished slides.

BURLEY, B. 2015. Extending Disney’s physically based BRDF with integrated subsurface scattering. *SIGGRAPH 2015 Physically Based Shading Course Notes*. (To appear).

D’EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4, 56:1–56:14.

D’EON, E., LUEBKE, D., AND ENDERTON, E. 2007. Efficient rendering of human skin. *Rendering Techniques (Proc. Eurographics Symposium on Rendering)*, 147–158.

D’EON, E. 2012. A better dipole. Tech. rep., <http://www-eugenedeon.com>.

D’EON, E. 2014. A dual-beam 3D searchlight BSSRDF. In *SIGGRAPH Talks*.

DONNER, C., AND JENSEN, H. W. 2007. Rendering translucent materials using photon diffusion. *Rendering Techniques (Proc. Eurographics Symposium on Rendering)*, 243–252.

DONNER, C., LAWRENCE, J., RAMAMOORTHY, R., HACHISUKA, T., JENSEN, H. W., AND NAYAR, S. 2009. An empirical BSSRDF model. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3, 30:1–30:10.

FRISVAD, J. R., HACHISUKA, T., AND KJELSDEN, T. K. 2014. Directional dipole model for subsurface scattering. *ACM Transactions on Graphics* 34, 1.

HABEL, R., CHRISTENSEN, P. H., AND JAROSZ, W. 2013. Photon beam diffusion: a hybrid Monte Carlo method for subsurface scattering. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 32, 4, 27–37.

JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21, 3, 576–581.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. *Computer Graphics (Proc. SIGGRAPH)* 35, 511–518.

JIMENEZ, J., ZSOLNAI, K., JARABO, A., FREUDE, C., AUZINGER, T., WU, X.-C., VON DER PAHLEN, J., WIMMER, M., AND GUTIERREZ, D. 2015. Separable subsurface scattering. *Computer Graphics Forum*. (To appear).

- KALOS, M. H., AND WHITLOCK, P. A. 1986. *Monte Carlo Methods*. John Wiley and Sons.
- KING, A., KULLA, C., CONTY, A., AND FAJARDO, M. 2013. BSSRDF importance sampling. In *SIGGRAPH Talks*.
- SCHLICK, C. 1994. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum 13*, 3, 233–246.
- WANG, L., JACQUES, S., AND ZHENG, L. 1995. MCML: Monte Carlo modeling of light transport in multi-layered tissues. *Computer Methods in Programs and Biomedicine*, 8, 313–371.
- YAN, L.-Q., ZHOU, Y., XU, K., AND WANG, R. 2012. Accurate translucent material rendering under spherical Gaussian lights. *Computer Graphics Forum (Proc. Pacific Graphics) 31*, 7, 2267–2276.