

# Photon Beam Diffusion: A Hybrid Monte Carlo Method for Subsurface Scattering


Ralf Habel<sup>1</sup>

Per H. Christensen<sup>2</sup>

Wojciech Jarosz<sup>1</sup>

<sup>1</sup>Disney Research Zürich

<sup>2</sup>Pixar Animation Studios



**Figure 1:** We solve the searchlight problem (a) using a Monte Carlo integration of dipoles distributed along either a normally incident (b) or oblique (c) beam. We can apply our hybrid approach to standard subsurface scattering (d) and also to more complex configurations using photon beams (e).

## Abstract

We present photon beam diffusion, an efficient numerical method for accurately rendering translucent materials. Our approach interprets incident light as a continuous beam of photons inside the material. Numerically integrating diffusion from such extended sources has long been assumed computationally prohibitive, leading to the ubiquitous single-depth dipole approximation and the recent analytic sum-of-Gaussians approach employed by Quantized Diffusion. In this paper, we show that numerical integration of the extended beam is not only feasible, but provides increased speed, flexibility, numerical stability, and ease of implementation, while retaining the benefits of previous approaches. We leverage the improved diffusion model, but propose an efficient and numerically stable Monte Carlo integration scheme that gives equivalent results using only 3–5 samples instead of 20–60 Gaussians as in previous work. Our method can account for finite and multi-layer materials, and additionally supports directional incident effects at surfaces. We also propose a novel diffuse exact single-scattering term which can be integrated in tandem with the multi-scattering approximation. Our numerical approach furthermore allows us to easily correct inaccuracies of the diffusion model and even combine it with more general Monte Carlo rendering algorithms. We provide practical details necessary for efficient implementation, and demonstrate the versatility of our technique by incorporating it on top of several rendering algorithms in both research and production rendering systems.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Ray Tracing; I.6.8 [Simulation and Modeling]: Simulation—Monte Carlo

## 1. Introduction

Many important materials exhibit noticeable translucency, including wax, marble, juices, the foods we eat, and our own skin. Accurate and efficient lighting simulation in such materials is a critical challenge in fields ranging from medical physics to realistic image synthesis in computer graphics. Light passes volumetrically through translucent materials before being scattered or absorbed, so they are poorly ap-

proximated by traditional surface shading models. In fact, subsurface scattering is just a simplified view of the more general problem of light transport in participating media like clouds, fog, smoke, and interstellar nebula.

Many rendering algorithms [PH10] numerically approximate light transport in participating media. Monte Carlo methods, such as variants of path tracing [Kaj86] or photon mapping [Jen96], are general since they can theoretically

handle arbitrary geometries and material properties. Unfortunately, these methods are typically slow to converge, especially when the medium is highly scattering and many bounces are needed to obtain converged results. Methods based on the diffusion approximation [Sta95], on the other hand, can be much more efficient in these situations, though they place strict assumptions on the medium being simulated.

In this paper, we present a new method for accurately rendering translucent materials called *photon beam diffusion* (PBD). Our approach interprets incident light as a continuous beam of photons, and expresses its propagation through the medium using enhancements to the state-of-the-art improved diffusion model [dl11]. Unlike previous work, we formulate our solution as a Monte Carlo integration problem. This provides increased flexibility, numerical stability, and speed. Furthermore, our formulation allows us to easily extend the accuracy and capabilities of the diffusion model and even combine it with more general Monte Carlo rendering algorithms.

### 1.1. Previous Work

**Diffusion-based Methods.** The diffusion equation can be solved using finite element methods [Sta95, AWB11, WZT\*08] which can handle heterogeneous materials and complex boundaries, but which require building the finite elements and solving the diffusion equation numerically, which is usually impractical for complex scenes.

Drawing on work from classical diffusion theory [FPW92], Jensen et al. [JMLH01] introduced the diffusion dipole to computer graphics and applied it to the rendering of translucent materials. Later, Donner and Jensen [DJ05] extended this to a multipole to accurately handle thin and multi-layer materials.

All of these methods use the classical diffusion theory approximation, which suffers from significant errors in near-source and high-absorption regimes. Recently, d'Eon and Irving [dl11] introduced, from other fields, a wealth of improvements to diffusion theory to address these problems, notably: 1) a modified diffusion-type equation from neutron transport [Gro56, Gro58] which explicitly decouples singly- and multiply-scattered light, and 2) a more accurate exitance calculation from optics [KP97]. D'Eon and Irving also promoted the use of an extended source term instead of approximating it as an impulse at a single depth like the dipole. Since no closed-form solution exists to the extended source integral, and available numerical approaches were expensive, they approximated the resulting diffusion profile as a sum of Gaussians. To avoid fitting the Gaussians [dLE07], d'Eon and Irving further exploited the fact that time-resolved or *quantized diffusion* (QD) results in a Gaussian distribution and used this as a mathematical basis for finding the Gaussian weights. Yan et al. [YZXW12] later extended the sum of Gaussians to handle oblique illumination from environment maps.

The improvements to diffusion theory, and the use of the extended source term, made QD accurate for a much wider range of translucent materials. Unfortunately, the *quantization* in QD has many practical drawbacks, including complicated and error-prone implementation, numerical instability, and high computational complexity due to the large number of Gaussians needed for accurate profiles.

In our work, we show that we can compute the same extended source integral numerically in a straightforward and very efficient way—eliminating the need for quantization and the sum of Gaussians altogether. We show that with just 3–5 numerical samples we can obtain a profile of equal quality to QD using 20–66 Gaussians, providing more flexibility and allowing us to easily approximate more general participating media transport than possible with QD by combining with other Monte Carlo rendering methods.

**Monte Carlo Methods.** Concurrently to the work on diffusion-based subsurface scattering methods, there have been significant advances in Monte Carlo rendering techniques for general participating media. We focus only on the most relevant previous work; a more comprehensive overview is provided by Pharr and Humphreys [PH10].

Variants of path tracing can handle arbitrary numbers of bounces, but at a significant cost and with high variance. Donner et al. [DLR\*09] used brute-force Monte Carlo particle tracing to tabulate an empirical BSSRDF model which handles oblique illumination and a broader range of albedo values than possible with diffusion. Kulla and Fajardo [KF12] recently proposed equiangular and decoupled importance sampling strategies to reduce noise when path tracing single-scattering. We employ equiangular sampling, but for computing multiple scattering with diffusion.

Monte Carlo algorithms that simulate multiple scattering more efficiently typically rely on some form of light path caching. Jensen and Christensen [JC98] generalized photon mapping to participating media by caching photons in volumes. Later, Jarosz et al. [JZJ08, JNSJ11] showed significant quality and performance improvements by considering the entire length of camera rays (the *beam radiance estimate*) or light paths (*photon beams*) when performing density estimation. This generalization relates to “track-length estimators” used in the neutron transport field for some time [Spa66]. Many-light methods [KHA\*12] turn photons into “virtual” light sources instead of using density estimation. The rendering pass then simulates an extra bounce of illumination by connecting these lights to shading points using visibility queries. Recently, Novak et al. [NNDJ12] developed an improved many-light method for media which, instead of turning photons into virtual point sources, turns photon beams into extended “virtual ray light” (VRL) sources—providing faster convergence and reduced singularities for multiple scattering. All of these Monte Carlo methods can simulate more general media than diffusion-based methods (accounting for

**Table 1:** Notation used throughout this paper

Symbol	Description	Units
$L(\vec{x}, \vec{\omega})$	Radiance at position $\vec{x}$ from direction $\vec{\omega}$	$[\text{W m}^{-2} \text{sr}^{-1}]$
$f_s(\vec{\omega}, \vec{\omega}')$	Phase function (normalized to $1/4\pi$ )	$[\text{sr}^{-1}]$
$g$	Average cosine of scattering	–
$\sigma_s$	Scattering coefficient	$[\text{m}^{-1}]$
$\sigma'_s$	Reduced scattering coefficient: $(1 - g)\sigma_s$	$[\text{m}^{-1}]$
$\sigma_a$	Absorption coefficient	$[\text{m}^{-1}]$
$\sigma_t$	Extinction coefficient: $\sigma_s + \sigma_a$	$[\text{m}^{-1}]$
$\sigma'_t$	Reduced extinction coefficient: $\sigma'_s + \sigma_a$	$[\text{m}^{-1}]$
$\eta$	Relative index of refraction	–
$\alpha$	Scattering albedo: $\sigma_s/\sigma_t$	–
$\alpha'$	Reduced scattering albedo: $\sigma'_s/\sigma'_t$	–
$\phi(\vec{x})$	Fluence	$[\text{W m}^{-2}]$
$\vec{E}(\vec{x})$	Vector flux (vector irradiance)	$[\text{W m}^{-2}]$
$Q(\vec{x})$	Source function	$[\text{W m}^{-3}]$
$D$	Diffusion coefficient	$[\text{m}]$
$\sigma_{tr}$	Transport coefficient: $\sqrt{\sigma_a/D}$	$[\text{m}^{-1}]$

heterogeneous media, anisotropic scattering and lighting distribution, and occlusions), but at generally higher cost, especially for highly scattering materials. As the name *photon beam diffusion* suggests, we take inspiration from this line of work by similarly considering a continuous “beam” of photons, but we simulate multiple scattering from this beam using diffusion instead of using density estimation or recursive Monte Carlo integration.

**Hybrid Methods.** A few hybrid methods have combined diffusion and Monte Carlo techniques. Li et al. [LPT05] proposed a hybrid approach which performs a few bounces of path tracing from the eye, and approximates subsequent bounces with the classical dipole. Donner and Jensen’s [DJ07] *photon diffusion* (PD) performs a photon tracing pass and then interprets the photons as diffusion sources. Like QD, this accounts for the extended incident beam of light, and further accounts for oblique incident illumination and internal blockers (approximately), but it forfeits accurate multilayered material support and uses the inferior classical diffusion model. We can likewise compute diffusion from photon sources, but operate on continuous photon beams [JNSJ11] instead of discrete photon points. This is similar in principle to many-light methods, but instead of computing exact one-bounce illumination from the photon sources, these methods compute approximate multi-bounce illumination using the diffusion approximation from photon points (for PD) or from photon beams (for our method, PBD).

## 1.2. Contributions

Our photon beam diffusion approach provides improvements across this wide variety of previous methods. Specifically:

- We provide a more efficient and numerically stable way to integrate the extended source function using Monte Carlo

importance sampling, eliminating the need for complex, numerically unstable quantization.

- We propose a new diffuse single-scattering term which exactly matches expensive Monte Carlo simulations.
- Our formulation of single and multiple scattering trivially allows for oblique illumination.
- We introduce an empirical correction factor which improves the accuracy of the modified diffusion model for near-source and oblique illumination regions.
- We demonstrate our method for both traditional subsurface scattering rendering and applied on top of Monte Carlo photon beam tracing, combining the complementary strengths of photon beams, virtual ray lights, photon diffusion, and the improved diffusion model.

## 2. Light Transport in Scattering Media

Light transport in a participating medium can be described by the *radiance transport equation* (RTE) [Cha60]:

$$(\vec{\omega} \cdot \vec{\nabla})L(\vec{x}, \vec{\omega}) = -\sigma_t L(\vec{x}, \vec{\omega}) + Q(\vec{x}, \vec{\omega}) + \sigma_s \int_{4\pi} L(\vec{x}, \vec{\omega}') f_s(\vec{\omega}, \vec{\omega}') d\vec{\omega}', \quad (1)$$

which states that the change in radiance  $L$  in direction  $\vec{\omega}$  is the sum of three terms: a decrease in radiance dictated by the extinction coefficient ( $\sigma_t = \sigma_s + \sigma_a$ ), and an increase due to the source function  $Q$  and the in-scattering integral on the second line. We summarize our notation in Table 1. In the most general form, simulating light transport in translucent materials requires solving the RTE (accounting for the scattering within the medium) with suitable boundary conditions imposed by the enclosing surfaces.

### 2.1. The BSSRDF and the Searchlight Problem

When rendering translucent materials, it is often convenient to re-formulate this problem in analogy to the local surface reflection integral. This results in an integral equation which computes the outgoing radiance,  $L_o$ , at position and direction  $(\vec{x}_o, \vec{\omega}_o)$  as a convolution of the incident illumination,  $L_i$ , and the BSSRDF,  $S$ , over all incident positions and directions  $(\vec{x}_i, \vec{\omega}_i)$ :

$$L_o(\vec{x}_o, \vec{\omega}_o) = \int_A \int_{2\pi} S(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o) L_i(\vec{x}_i, \vec{\omega}_i) (\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i dA(\vec{x}_i). \quad (2)$$

For efficiency,  $S$  is often decomposed into reduced radiance, single scattering, and multi-scattering terms,  $S = S^{(0)} + S^{(1)} + S_d$ , so that each can be handled by specialized algorithms. Most previous techniques in graphics have approximated  $S^{(1)}$  with the refractive single-scattering approximation proposed by Jensen et al. [JMLH01]. Simulating the multi-scattering term  $S_d$  accurately is expensive, so techniques in fields such as medical physics, astrophysics, and recently computer graphics have relied on simplifications to this general problem based on approximate solutions to the so-called “searchlight problem.”

The searchlight problem, Figure 1(a), considers a simplified setting where a focused beam of light is incident at the origin on a semi-infinite planar homogeneous medium. Photons refract through a Fresnel boundary and travel in the downward direction until they are scattered by the medium, and ultimately get absorbed or escape the material. The distribution of photons exiting the upper boundary forms a spatial reflectance profile  $R_d(\vec{x})$ , which is radially symmetric (1D) for normally-incident light,  $R_d(\vec{x}) = R_d(\|\vec{x}\|)$ . General solutions to the searchlight problem typically rely on brute-force Monte Carlo particle tracing for validation, though exact solutions do exist for special cases.

Most methods in graphics simplify  $S_d$  as a product of a typically 1D, diffuse reflectance profile  $R_d$  and a directional, energy-preserving Fresnel reshaping term [JMLH01, dl11]:

$$S_d(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o) = \frac{1}{\pi} F_t(\vec{x}_i, \vec{\omega}_i) R_d(\vec{x}_o - \vec{x}_i) \frac{F_t(\vec{x}_o, \vec{\omega}_o)}{4C_\phi(1/\eta)}, \quad (3)$$

where the reflectance profile is now centered at the incident light position,  $\vec{x}_i$ , and  $C_\phi$  is a constant needed for normalization. To evaluate this expression efficiently, most methods have relied on diffusion theory to obtain an analytic approximation for  $R_d$ .

For a full review of diffusion *theory* and the improvements which were previously described in tandem with the proposed *solution* using quantized diffusion and Gaussians [dl11], please refer to our supplemental technical report [HCJ13]. In the following, we only present a short description resulting in the integral from the improved diffusion model that needs to be solved. We then show our solution using a Monte Carlo approach, which is independent of the underlying model used. Furthermore, we show how the same integration scheme can be used to efficiently compute the *exact* single-scattering from the searchlight problem and use it to define a new diffuse single-scattering term  $S^{(1)}$ .

## 2.2. Improved Diffusion Model and Extended Source

The reflectance profile  $R_d$  can be obtained by considering an exponentially-decreasing source beam within the material and expressing its scattering and propagation using the diffusion approximation. This results in the following integral:

$$R_d(\vec{x}, \vec{\omega}) = \int_0^\infty r(\vec{x}, \vec{x}_r(t)) Q(t) dt, \quad (4)$$

where  $Q(t) = \alpha' \sigma'_t e^{-\sigma'_t t}$  is the extended source function, and  $r(\vec{x}, \vec{x}_r(t)) = R_\phi^d(\vec{x}, t) + R_E^d(\vec{x}, t)$  is the radiant intensity on the surface due to a dipole along the beam at  $\vec{x}_r(t) = t\vec{\omega}$  (see Figure 1(b)). The radiant intensity depends on fluence and vector irradiance terms, defined as:

$$R_\phi^d(\vec{x}, t) = C_\phi \frac{\alpha'}{4\pi D} \left( \frac{e^{-\sigma_{tr} d_r(t)}}{d_r(t)} - \frac{e^{-\sigma_{tr} d_v(t)}}{d_v(t)} \right), \quad (5)$$

$$R_E^d(\vec{x}, t) = C_E \frac{\alpha'}{4\pi} \left[ \frac{z_r(t) (1 + \sigma_{tr} d_r(t)) e^{-\sigma_{tr} d_r(t)}}{d_r^3(t)} + \frac{(z_r(t) + 2z_b) (1 + \sigma_{tr} d_v(t)) e^{-\sigma_{tr} d_v(t)}}{d_v^3(t)} \right], \quad (6)$$

where we use the shorthand  $d_r(t) = d(\vec{x}, \vec{x}_r(t))$  and analogously for  $d_v$ , and where  $z_r(t) = \vec{x}_r(t) \cdot \vec{n}$  is the depth of the real source.


## 2.3. Quantized Diffusion

D'Eon and Irving noted that Equation (4) has no closed-form solution, and proposed to approximate it using a sum of Gaussians:

$$R_d(\vec{x}) \approx \alpha' \sum_{i=0}^{k-1} (w_R(i) w_i) G_{2D}(v_i, x), \quad (7)$$

where  $x = \|\vec{x}\|$ ,  $(w_R(i) w_i)$  are the weights, and  $v_i$  are the variances of normalized 2D Gaussians  $G_{2D}$ . The equations necessary to obtain the weights and variances of the Gaussians are themselves complex summations of integrals depending on further weights  $w_{\phi R}(v, i)$ ,  $w_{ER}(v, i)$ , which we omit here.

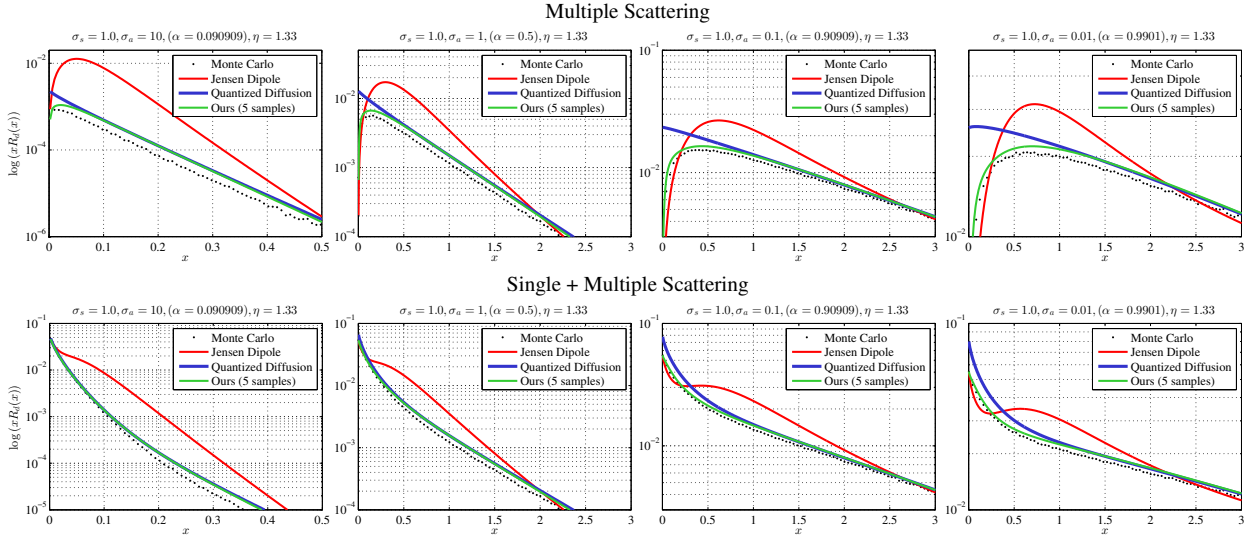
We found that there are several practical downsides of solving Equation (4) in this manner. Evaluating Equation (7) is non-trivial, relatively expensive, and prone to error since the formulas provided for the weights are highly susceptible to numerical instabilities—specifically, when incorporating the extinction of the source beam into the Gaussians weights. Figure 2 shows a plot of the weights where the instability occurs.



**Figure 2:** Gaussian weights computed as described by d'Eon and Irving [dl11] with double precision (blue) show strong instabilities compared to the expected behavior (green).

Though this problem can be circumvented by deriving functional approximations (kindly provided to us by the authors [dl11]), this adds to the already high complexity of the method. Moreover, Gaussian functions are actually a poor representation for diffusion profiles, and therefore many Gaussians are needed for an accurate reconstruction. Due to the complexity of the calculations and the need for many Gaussians, a quantized diffusion profile is roughly an order of magnitude more expensive to evaluate than the standard dipole.

On the other hand, the profiles resulting from quantized diffusion are significantly more accurate than the standard dipole—see Figure 3. The downsides are due to the



**Figure 3:** Comparison of the standard dipole, QD, and our method (with  $\kappa$  correction) validated against Monte Carlo simulation for multiple scattering only (top row) and combined with single scattering (bottom row). Our approach matches the quantized diffusion profile, but uses only 5 samples compared to 66 Gaussians. By incorporating our  $\kappa$  correction factor, we are able to surpass the fidelity of QD, avoiding divergence near the source and matching Monte Carlo much more faithfully.

quantization procedure itself (which is orthogonal to the other improvements to diffusion imported to graphics by d’Eon and Irving—see also the recent work by d’Eon and Hery [d’E12, Her12]). This motivated us to leverage the improved model but solve Equation (4) in a different way.

### 3. Method

The core of our method is to numerically compute the extended source integral (4) using Monte Carlo integration with importance sampling:

$$R_d(\vec{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{i=1}^N \frac{f(\vec{x}, \vec{\omega}, t_i)}{\text{pdf}(t_i | \vec{x}, \vec{\omega})}, \quad (8)$$

where  $f(\vec{x}, \vec{\omega}, t_i) = r(\vec{x}, \vec{x}_r(t_i)) Q(t_i)$  is the integrand from (4) and  $\text{pdf}(t_i | \vec{x}, \vec{\omega})$  is the PDF of choosing  $t_i$  given the shading location  $\vec{x}$  and the beam direction  $\vec{\omega}$ .

Computing the integral in this way provides many practical advantages. Firstly, it allows us to define our set of “effective basis functions” (the summands of the above equation) by carefully choosing the PDF. This allows us to obtain accurate results with as little as  $N = 3.5$  samples while supporting asymmetric profiles from oblique incident illumination (see Figure 1(b,c)), finite slabs, and multilayer materials. We also introduce a correction term which improves the accuracy of the underlying diffusion model near the incident beam, especially for highly scattering materials. This improvement is not possible to incorporate into QD, but fits trivially within our more flexible integration scheme. Finally, our approach al-

lows us to efficiently integrate a new diffuse single-scattering term motivated by the searchlight problem.

#### 3.1. Importance Sampling

To evaluate Equation (8), we can use any number of importance sampling strategies to distribute the numerical samples.

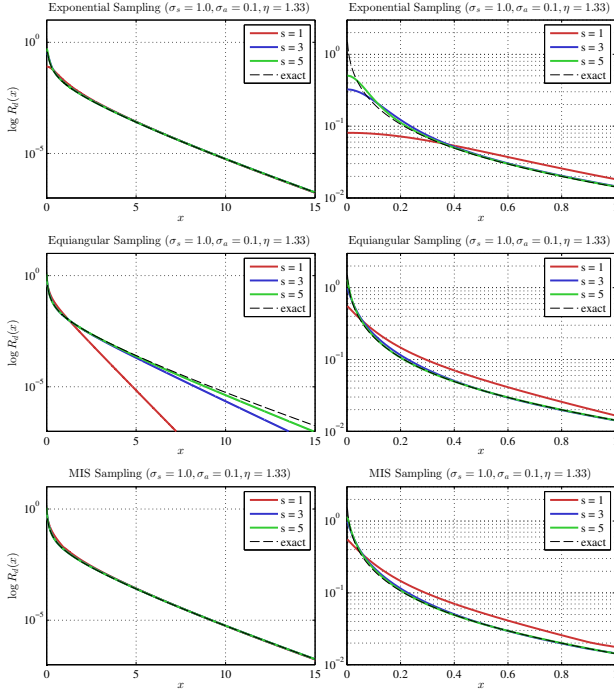
**Exponential Sampling.** The traditional approach—used by path tracing, photon mapping, and photon diffusion for choosing propagation distances in homogeneous media—is exponentially-decreasing sampling with a PDF proportional to the source term  $Q(t)$ :

$$t_i = -\frac{\log(1 - \xi_i)}{\sigma'_i} \quad \text{with} \quad \text{pdf}_{\text{exp}}(t_i) = \sigma'_i e^{-\sigma'_i t_i}, \quad (9)$$

where  $\xi_i \in [0, 1)$  is a uniform random number. This PDF only depends on  $t$ , and is independent of  $\vec{x}$  and  $\vec{\omega}$ . This effectively places dipole sources at different depths and weights them according to the inverse exponential PDF, canceling out the exponential variation in  $Q(t)$ .

**Equiangular Sampling.** A different strategy is, at each  $\vec{x}$ , to distribute the samples uniformly in the angular domain subtended by the incident beam (Figure 1(b)). This corresponds to the equiangular sampling method proposed by Kulla and Fajardo [KF12] and also used by Novák et al. [NNDJ12]. This gives:

$$t_i = -h \tan(\theta_i) \quad \text{with} \quad \text{pdf}_{\text{equ}}(t_i | \vec{x}, \vec{\omega}) = \frac{h}{(\theta_b - \theta_a)(h^2 + t_i^2)} \quad (10)$$



**Figure 4:** The profiles (left) and close-up of the peak (right) of the sampling methods for different numbers of samples. Exponential sampling always introduces some smoothing near  $x = 0$  while equiangular sampling cannot match the tail behavior. Our combined MIS strategy (bottom) matches both. Profiles for other absorption values are available in the supplemental document.

where  $h$  is the distance between  $\vec{x}$  and the beam  $\vec{\omega}$ ,  $\theta_a$  and  $\theta_b$  are the start and end angles of the integration in the angular domain, and  $\theta_i = \text{lerp}(\theta_a, \theta_b, \xi_i)$ .

**Discussion & Comparison.** Standard Monte Carlo integration uses uniform random numbers for the  $\xi_i$ 's. To avoid per-pixel noise in our profiles, we instead use a deterministic regular sequence  $\xi_i = \frac{i-0.5}{N}$ . This effectively creates a custom quadrature rule which is defined by the positions and weights of the importance sampling method used, directly using information from the integral through the scattering parameters. We found that this produces superior results and achieves high accuracy with a very low number (3–5) of samples.

Figure 4 (top and middle) compares the convergence behavior of exponential and equiangular sampling. In the exponential case, the tail of the profile is guaranteed to converge when taking just 1 sample. On the other hand, equiangular sampling represents the peak well even with 1 sample, whereas exponential sampling underestimates the peak.

**Combining Strategies.** Given two sampling strategies that perform well in different situations the standard solution is to use multiple importance sampling [VG95] (MIS) to combine

the strategies. Using the standard balance heuristic would draw samples from both strategies across the entire  $\vec{x}$  domain. This is wasteful and would not guarantee proper asymptotic tail behavior. Luckily, we always know which strategy works better in which region (tail vs. peak) and use this knowledge to develop a modified MIS weighting heuristic, tailored to our specific problem. Our estimator becomes:

$$R_d \approx \frac{1}{N} \sum_{i=1}^N \frac{f(\vec{x}, \vec{\omega}, t_i) w_{\text{exp}}(t_i, x)}{w_{\text{exp}}(t_i, x) \text{pdf}_{\text{exp}}(t_i) + w_{\text{eq}}(t_i, x) \text{pdf}_{\text{eq}}(t_i | \vec{x}, \vec{\omega})} + \frac{1}{N} \sum_{j=1}^N \frac{f(\vec{x}, \vec{\omega}, t_j) w_{\text{eq}}(t_j, x)}{w_{\text{exp}}(t_j, x) \text{pdf}_{\text{exp}}(t_j) + w_{\text{eq}}(t_j, x) \text{pdf}_{\text{eq}}(t_j | \vec{x}, \vec{\omega})}, \quad (11)$$

where  $w_{\text{exp}}(t_i, x) = \text{clamp}((x-a)/(b-a), 0, 1)$  and  $w_{\text{eq}}(t_i, x) = 1 - w_{\text{exp}}(t_i, x)$ . Our weighting strategy dedicates samples entirely to the superior PDF except in a small transition zone ( $x \in [a, b]$ ) where we smoothly weight the two strategies. We center the transition zone around one diffuse mean free path, with  $a = 0.9/\sigma'_i$  and  $b = 1.1/\sigma'_i$ . This approach can be interpreted as deriving domain-specific “representativity” for the two strategies used in the MIS estimator [BPBP11].

In Figure 4 (bottom) we show how our tailored MIS strategy, with just 3 samples, matches the entire profile, including the peak and the asymptotic tail behavior.


### 3.2. Correction Factor

Donner and Jensen [DJ07] noticed an overestimation of exitance near the incident light source when using the classic dipole and proposed an empirical correction factor  $\kappa(t) = 1 - e^{-\sigma t}$  that diminishes near-surface regions of  $Q(t)$ . D’Eon and Irving [dl11] instead used the improved model with Kienle-Patterson exitance calculations and showed that the resulting profiles were consistently more accurate. However, we found that this improved model also suffers from overestimation for near-surface sources. Figure 3 (top) shows that QD’s multiple scattering in fact diverges as  $\vec{x}$  approaches 0, though the finite number of Gaussians avoids an actual singularity at  $\vec{x} = 0$ . To account for this, we introduce an empirical correction factor inspired by Donner and Jensen’s, but tailored for the improved diffusion model:

$$\kappa(\vec{x}, t) = 1 - e^{-2\sigma_i(d(\vec{x}, \vec{x}_r(t))+t)}. \quad (12)$$

In Figure 3 we compare our approach to both the classic dipole and to QD (using 66 Gaussians) for a semi-infinite medium with fixed  $\sigma_s = 1$  and three orders of magnitude of absorption levels. We also include a ground truth computed using our own brute-force Monte Carlo particle tracing code, which we first validated against the well-tested MCML program [WJZ95].

While QD strongly overestimates at  $\vec{x} \rightarrow 0$ , our profile with the correction factor matches the shape of the Monte Carlo reference profile well over the entire  $\vec{x}$  domain and absorption parameter range. It is important to point out that since our correction factor depends not only on  $t$ , but also




**Figure 5:** Our method matches ground truth multi-scattering (top) better than previous approaches, including QD (no correction factor), better dipole, and classic dipole. Our diffuse single scattering (bottom) computed with just 5 samples, matches Monte Carlo reference faithfully.

on the shading location  $\vec{x}$  (indirectly through its dependence on  $d_r = d(\vec{x}, \vec{x}_r(t))$ ), it is impossible to incorporate it in the sum-of-Gaussians approach used by QD. Our Monte Carlo integration strategy, however, makes incorporating this correction factor trivial.

Figure 5 shows a comparison of the different methods. The scene consists of a semi-infinite medium with the surface illuminated by three thin, bright stripes of light. We use  $\sigma'_s = 1, \sigma_a = (0.01, 0.1, 1), \eta = 1, g = 0$  for all validation results to show low, mid, and high absorptive media in a single color image. We include all results as HDR images in the supplemental material and encourage the reader to view them at different exposures.

### 3.3. Finite Slabs and Multilayer Materials

Handling finite slabs with our numerical integration scheme is straightforward. We simply treat each sample (summand of Equation (8)) as a dipole and mirror it at the extrapolated boundary to perform a multipole extension for a finite slab as in Donner and Jensen [DJ05]. This has the same effect as the mirroring of Gaussians in QD. The equiangular sampling in our MIS strategy avoids placing dipoles directly on the reflective as well as transmissive side of the slab when close to  $\vec{x} = 0$  which allows for very thin slabs without any drawbacks. Furthermore, our  $\kappa$  correction factor significantly improves reflectance as well as transmittance compared to the improved diffusion model for finite slabs as shown in Figure 6 below.



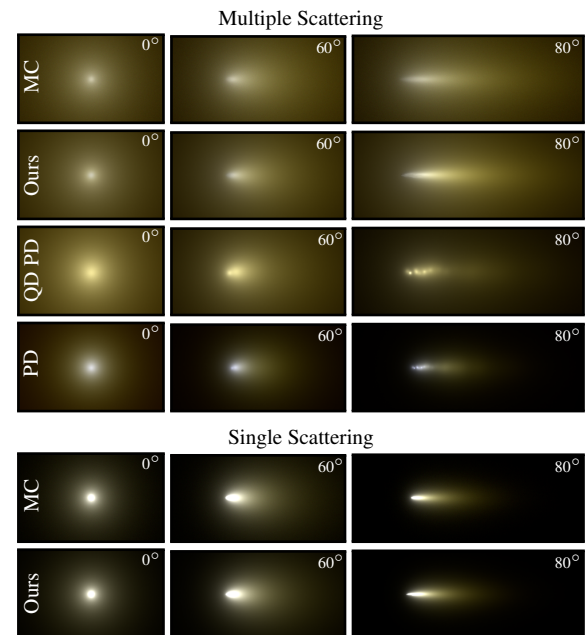
**Figure 6:** Reflectance (left) and transmittance (right) of a finite slab with thickness  $d = 0.1/\sigma_s$ .

Multilayer materials can be handled in the same way as Donner and Jensen. Since Gaussians naturally support convolution required for multi-layer materials, QD admittedly holds an advantage here over our method or Donner and Jensen's original multipole formulation.

### 3.4. Oblique Incident Light Angles

Similar to Yan et al. [YZXW12], who formulated QD for oblique angles by offsetting the centers of the Gaussians, we can apply our approach to oblique light angles. Due to the explicit integration formulation, Equation (8) trivially allows us to choose any  $\vec{\omega}$  that is not orthogonal to the surface normal without any modifications. We simply compute locations along the refracted incident beam during numerical integration as shown in Figure 1(c).

In Figure 7 we evaluate the accuracy of our method for oblique angles on the same medium as in Figure 5. We compare to Monte Carlo for validation and we also compare to PD with photons deposited using an exponential distribution [DJ07] for both the original PD formulation as well as PD with the improved diffusion model. Steep incident angles pose problems due to both the integration/sampling method



**Figure 7:** An incident light beam (with radius  $0.2/\sigma_s$ ) illuminates a semi-infinite medium at angles  $0^\circ, 60^\circ$  and  $80^\circ$  degrees from normal incidence. We compare Monte Carlo reference (top), our multi-scattering (middle) using 50 beams, and photon diffusion [DJ07] with the improved model (QD PD) and classic model (PD) using an equivalent number of samples ( $50 \times 5 = 250$  photons). The bottom compares Monte Carlo single-scattering reference to our solution with 5 samples. (Parameters:  $\sigma'_s = 1, \sigma_a = (0.01, 0.1, 1), \eta = 1, g = 0$ .)

and approximation in the diffusion model. Our results match Monte Carlo references accurately at a wide variety of incident angles and do not break down even at extreme grazing angles (see supplemental material). The improvement in overall intensity and color compared to photon diffusion is due to the improved diffusion model and our  $\kappa$  correction factor, which becomes increasingly important for avoiding overestimation at steep incident angles. Also, even with the same number of total dipole evaluations, our method benefits from better sample placement which adapts the “photon” locations to each shading location.

### 3.5. Diffuse Single Scattering

Single scattering is an important part of translucent material appearance, especially near the incident light and at high absorption. Jensen et al. [JMLH01] derived an approximate single-scattering term assuming smooth Fresnel boundary conditions. Unfortunately, this usually results in a very dim contribution, and is therefore often omitted from production implementations. However, since the improved model separates out single scattering we need to account for it explicitly.

Our goal is to include the same diffuse single-scattering term computed by brute-force Monte Carlo particle tracing simulations such as MCML [WJZ95]. This type of single-scattering term has typically been used only in validation graphs, but not incorporated in any practical rendering algorithm since using brute-force particle tracing during rendering would be prohibitively expensive. However, since we are explicitly integrating along the incident beam using Monte Carlo integration, we can efficiently and trivially include single scattering in the same way as the multi-scattering.

In fact, the only difference is that the Green’s function (the  $r(\vec{x}, \vec{x}_r(t))$  term in the integrand of Equation (4)) needs to change to a single-scattering response:

$$r^{(1)}(\vec{x}, \vec{x}_r(t)) = \frac{f_s(\vec{\omega} \cdot \vec{\omega}_{\vec{x}, \vec{x}_r}) e^{-\sigma_t(d_r(t))} F_t(\theta_o, \eta) F_t(\theta_i, 1/\eta) \cos \theta_o}{d_r^2(t)} \quad (13)$$

where  $\vec{\omega}_{\vec{x}, \vec{x}_r}$  and  $d_r(t)$  are the direction and distance from  $\vec{x}_r(t)$  to  $\vec{x}$  respectively,  $\theta_o = \arccos(\vec{\omega}_{\vec{x}, \vec{x}_r} \cdot \vec{n})$ ,  $\theta_i = \arccos(-\vec{\omega} \cdot \vec{n})$ , and  $F_t$  and  $f_s$  are the transmissive Fresnel and phase function terms. Figure 1(b,c) illustrates the configuration.

We can then compute the single-scattering profile,  $R^{(1)}(\vec{x}, \vec{\omega})$ , using the same Monte Carlo estimator (8), but with  $f(\vec{x}, \vec{\omega}, t_i) = r^{(1)}(\vec{x}, \vec{x}_r(t_i)) Q(t_i)$ . The single-scattering portion of the BSSRDF,  $S^{(1)}$ , can then be defined in direct analogy to  $S_d$  (3), but using  $R^{(1)}$  instead of  $R_d$ . We can compute the two profiles simultaneously (using the same points or rays), adding minimal overhead to the calculations. However, if the surface separates volumes with different indices of refraction, we can use a separate sample sequence for single-scattering, starting the integration not from the incident light angle, but from the critical angle since contributions below the criti-

cal angle are zero. This optimization guarantees a smooth single-scattering profile regardless of the number of samples.

In contrast to multi-scattering using diffusion, the single-scattering integral is exact under the assumed planar configuration and boundary conditions (Figure 1). In particular, the phase function and Fresnel terms are accounted for explicitly instead of approximately using similarity theory and the reflection parameter. Also, since we calculate a single bounce, no negative sources (dipoles) appear in the integrand. It is worth pointing out that this equation is mathematically similar to computing the diffuse illumination on a surface due to a VRL [NNDJ12] but with the addition of Fresnel terms. Unlike the refractive single-scattering term of Jensen et al. [JMLH01], which is highly directional and dim, our diffuse single scattering is a significant part of the overall radiant exitance profile (see Figures 3, 5, and 7).

Compared to single scattering at a Fresnel boundary as described by Jensen et al. [JMLH01], we assume the exitant boundary is diffuse—though not Lambertian, since we reshape the exitant radiance by a Fresnel term as done with multi-scattering. Jensen et al. on the other hand approximate the length of the incident beam inside the medium with a shadow ray that ignores refraction. Both of these are approximations. Jensen et al.’s single scattering approach is likely more amendable to media such as marble or milk with a smooth Fresnel boundary while our model caters more to skin or other materials with rough boundaries. A full investigation of the relative strengths of these two approaches would be an interesting avenue for future work.

## 4. Implementation and More Results

In addition to the scientific results for validation in the previous sections, we additionally evaluated photon beam diffusion on practical scenes. To demonstrate the versatility of our approach, we incorporated it into several rendering systems and light transport algorithms, including a point-based subsurface scattering implementation in Mitsuba [Jak10], point-based and ray-traced subsurface scattering in PRMan, and on top of a progressive photon beam tracer [JNT\*11] to simulate oblique illumination.

### 4.1. Subsurface Scattering with PBD Profiles

Our Mitsuba point-based implementation first distributes surface points using point repulsion and writes them in a point cloud file. It then constructs an octree of the points, and hierarchically traverses the octree [JB02]. Figure 8 shows a direct comparison between volumetric path tracing, photon beam diffusion, and dipole diffusion rendered in Mitsuba.

PRMan’s implementation of subsurface scattering encompasses both point-based and ray-traced approaches. For the point-based approach, the object surfaces are first tessellated into micropolygons and shaded, and the shaded micropolygons are written as points in a point cloud file. PRMan then





**Figure 8:** Comparison between volumetric path-tracing (top left), our method (top right), our method with multiple scattering only (bottom left), and dipole diffusion (bottom right). The dipole produces color shifts and over-smoothing due to an inaccurate profile. (Parameters:  $\sigma'_s = 1, \sigma_a = (0.01, 0.1, 1), \eta = 1, g = 0$ .)

constructs an octree over the points, and traverses the it hierarchically during rendering [JB02]. Ray-traced subsurface scattering at a shading point is computed by selecting a point inside the object (below the shading point) and shooting rays in all directions. We compute the area, incident illumination, and subsurface contribution for each ray hit point. The incident illumination can be cached in a multiresolution irradiance cache for efficiency [CHS\*12].

The interface and parameters of our method are the same as for the classic dipole, though the different methods cause a few practical issues that we address in the following.

#### 4.1.1. Efficient Profile Evaluation

Computing high-quality subsurface scattering at a point requires evaluating the BSSRDF at a few hundred distances—either for subsurface ray hit points or for point and cluster positions. Evaluating the BSSRDF hundreds of times for each shading point is relatively slow.

For efficiency, d’Eon and Irving [dl11] used binning to reduce the number of BSSRDF evaluations: they deposit cache point energy into discrete radial bins during hierarchical traversal, and then multiply each bin by the integrated profiles to accumulate the final energy. This is relatively efficient and accurate since their Gaussian profile can be analytically integrated within each bin. We take a dual approach by tabulating the profile itself, to achieve a similar effect without requiring bin integration.

Like d’Eon and Irving, for each shade point, before traversing the point hierarchy or shooting rays, we choose an exponentially-varying set of distances  $x = 0, x_1, x_1\Delta, x_1\Delta^2, x_1\Delta^3, \dots$ , with a step size  $\Delta$  of 1.2 and  $x_1 \approx 0.01/\sigma_r$ . But instead of using these distances as discrete bins, we evaluate and cache the reflectance profile at these distances. During

ray tracing or octree traversal, we use the ray length or the distance to the octree point or cluster and interpolate between the two nearest cached profile values.

Though caching alleviates evaluation time impact on the overall rendering time, efficient generation of profiles is important if the scattering parameters change over the surface and a new profile needs to be generated for each shading point. With the QD weight calculation and the evaluation of the Gaussians, we achieve a speedup factor of  $3.1\times$  for a full profile, comparing 5 photon beam diffusion samples with 44 Gaussians. In the less common case where only one sample of a profile is needed, we achieve a speedup of  $21.6\times$ .

#### 4.1.2. Parameter Conversion

For practical applications, it is more user-friendly to parameterize the BSSRDF by diffuse albedo  $R_d$  (diffuse reflection coefficient) and diffuse mean free path length ( $1/\sigma'_d$ ) than by  $\sigma'_s$  and  $\sigma_a$ . For dipole diffusion, Jensen and Buhler [JB02] derived an expression for  $R_d$  given  $\alpha'$  and noted that although the expression isn’t analytically invertible it is monotonically increasing so the inversion can be done using secant root finding. We have developed a more efficient table-based method for this inversion.


For a given index of refraction, we first generate a table of  $R_d$  values for values of  $\alpha' \in [0, 1)$ . Next, the table is inverted, resulting in a table of  $\alpha'$  values for values of  $R_d$ . Interpolating between two table entries, we can determine  $\alpha'$  while avoiding expensive profile evaluations. Together with the diffuse mean free path, we then get  $\sigma'_s$  and  $\sigma_a$ .

If we had used secant root finding, the  $R_d$  evaluations required for photon beam diffusion would be more expensive than for dipole diffusion: for dipole diffusion, computing  $R_d$  given  $\alpha'$  only requires evaluating an analytic formula, while for photon beam diffusion we need to perform a numerical integration. But fortunately the same table inversion method can be used for both diffusion models, making their parameter conversion times nearly identical. Thanks to these table lookups and efficient profile evaluations, rendering with photon beam diffusion and single-scattering typically takes only about 15% longer than using the classic dipole.


Figure 9 shows a human head with ray-traced dipole diffusion compared against photon beam diffusion with and without single-scattering. Subsurface scattering with photon beam diffusion is much sharper than dipole diffusion around skin pores and bumps on the lips, and single scattering adds further sharpness. A turntable animation of the head and a complex scene rendered with photon beam diffusion can be found in the supplemental material.

## 4.2. Applying PBD to Photon Beams

In addition to standard subsurface scattering rendering which convolves the diffusion profile with the direction-less surface



**Figure 10:** Photon beam diffusion shows a faster multiple scattering convergence than PD due to the equiangular/exponential MIS sampling, especially at oblique light angles (see insets). The left image shows our multiple scattering (MS) combined with single scattering (SS) calculated using photon beams.



**Figure 9:** The classic dipole model (left) can lead to an overly waxy appearance for human skin. Our PBD model (middle) produces sharper profiles, and additionally incorporating our diffuse single scattering term (right) preserves pore details even better. Render times:  $\sim 7$  minutes on a multi-core PC.

irradiance, we have applied our improved profiles and numerical integration scheme on top of a photon beam tracer to simulate diffusion from directional, oblique illumination. Our implementation operates much like PD [DJ07], but we store continuous, directional photon beams entering the surface instead of discrete photon points inside the medium (see Figure 1(e) and compare to 1(d)). We use the stored photon beams as extended beam sources and evaluate the multiple scattering using our integration scheme and the improved diffusion model. Applying photon beam diffusion to oblique illumination on arbitrary geometries requires more careful handling of the boundary conditions. We blend between dipole, multipole, and quadrupole boundary conditions as proposed by Donner and Jensen [DJ07]. Our Monte Carlo integration scheme (8) makes the link between these two explicit: each Monte Carlo sample of the beam integral can be viewed as a weighted photon at  $\vec{x}_r(t_i)$  inside the medium (a monopole), which we can mirror exactly as in PD to compute the three boundary conditions. By storing photon beams instead of discrete points during photon tracing, we can compute different

photon locations along the beams for each shading location using equiangular sampling. Figure 10 shows a comparison of the ground truth, photon diffusion, and our method, showing a faster convergence with less noise with our equiangular diffusion sampling.

## 5. Conclusion and Future Work

We have presented a hybrid method—photon beam diffusion—that combines Monte Carlo integration with the diffusion approximation. We can obtain the same improved subsurface scattering results as quantized diffusion, but our method is simpler, faster, more accurate, and numerically stable. On the other hand, certain features such as multi-layered materials are faster and more elegant in QD than with our explicit integration. Our method can be used with point-based and ray-traced subsurface scattering, and is only a few percent slower than the classic dipole diffusion model. We integrated it in both research and production renderers to validate the method. We showed that photon beam diffusion can be combined with more general rendering algorithms like photon beams, minimizing artifacts associated with strongly oblique angles.

Our versatile numerical integration scheme opens up the doors for several improvements. It is theoretically possible to account for more general rough dielectric refraction distributions for both multi- and single-scattering, by incorporating them directly into our numerical integration process. This would provide a principled way to handle more complex BSDFs in place of our current diffuse distribution at the boundary. We also believe that a full combination of bidirectional path tracing and diffusion approximation is possible, though the connection to the path vertices requires a more fundamental reformulation of diffusion.

**Acknowledgements.** We would like to thank Eugene d’Eon for explaining and clarifying QD and for providing the functional approximations. We thank Marios Papas for helpful discussions. We thank Maurizio Nitti for the models shown in the supplemental material and Chris Harvey for early testing and for rendering the head images and animation. The head data is courtesy of Infinite Realities and available via Creative Commons.

## References

- [AWB11] ARBREE A., WALTER B., BALA K.: Heterogeneous subsurface scattering using the finite element method. *IEEE Transactions on Visualization and Computer Graphics* 17, 7 (2011), 956–969. 2
- [Cha60] CHANDRASEKAR S.: *Radiative Transfer*. Dover, 1960. 3
- [CHS\*12] CHRISTENSEN P. H., HARKER G., SHADE J., SCHUBERT B., BATALI D.: *Multiresolution radiosity caching for efficient preview and final quality global illumination in movies*. Tech. Rep. 12-06, Pixar, 2012. 9
- [d'E12] D'EON E.: *A better dipole*. Tech. rep., <http://www.eugenedeon.com>, 2012. 5
- [dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4 (2011), 56:1–56:14. 2, 4, 6, 9
- [DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 24, 3 (2005), 1032–1039. 2, 7
- [DJ07] DONNER C., JENSEN H. W.: Rendering translucent materials using photon diffusion. *Rendering Techniques (Proc. Eurographics Symposium on Rendering)* (2007), 243–252. 3, 6, 7, 10
- [dLE07] D'EON E., LUEBKE D., ENDERTON E.: Efficient rendering of human skin. *Rendering Techniques (Proc. Eurographics Symposium on Rendering)* (2007), 147–158. 2
- [DLR\*09] DONNER C., LAWRENCE J., RAMAMOORTHI R., HACHISUKA T., JENSEN H. W., NAYAR S.: An empirical bssrdf model. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (July 2009), 30:1–30:10. 2
- [FPW92] FARRELL T. J., PATTERSON M. S., WILSON B.: A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Medical Physics* 19, 4 (1992), 879–888. 2
- [Gro56] GROSJEAN C. C.: A high accuracy approximation for solving multiple scattering problems in infinite homogeneous media. *Il Nuovo Cimento (1955-1965)* 3 (1956), 1262–1275. 2
- [Gro58] GROSJEAN C. C.: Multiple isotropic scattering in convex homogeneous media bounded by vacuum. In *Proc. Second International Conference on the Peaceful Uses of Atomic Energy* (1958), vol. 413. 2
- [HCJ13] HABEL R., CHRISTENSEN P. H., JAROSZ W.: *Classic and modified diffusion theory for subsurface scattering*. Tech. rep., Disney Research Zürich, 2013. 4
- [Her12] HERY C.: *Texture mapping for the better dipole model*. Tech. Rep. 12-11, Pixar, 2012. 5
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 8
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21, 3 (2002), 576–581. 8, 9
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. *Computer Graphics (Proc. SIGGRAPH)* 32 (1998), 311–320. 2
- [Jen96] JENSEN H. W.: Global illumination using photon maps. *Rendering Techniques (Proc. Eurographics Workshop on Rendering)* (1996), 21–30. 1
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. *Computer Graphics (Proc. SIGGRAPH)* 35 (2001), 511–518. 2, 3, 4, 8
- [JNSJ11] JAROSZ W., NOWROUZEZAHRAI D., SADEGHI I., JENSEN H. W.: A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* 30, 1 (2011), 5:1–5:19. 2, 3
- [JNT\*11] JAROSZ W., NOWROUZEZAHRAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6 (2011), 181:1–181:12. 8
- [JZJ08] JAROSZ W., ZWICKER M., JENSEN H. W.: The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum (Proc. Eurographics)* 27, 2 (2008), 557–566. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (Proc. SIGGRAPH)* 20 (1986), 143–150. 1
- [KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 31, 4 (2012), 1519–1528. 2, 5
- [KHA\*12] KRIVÁNEK J., HAŠAN M., ARBREE A., DACHSBACHER C., KELLER A., WALTER B.: Optimizing realistic rendering with many-light methods. In *SIGGRAPH 2012 Courses* (2012). 2
- [KP97] KIENLE A., PATTERSON M. S.: Improved solutions of the steady-state and the time-resolved diffusion equations for reflectance from a semi-infinite turbid medium. *Journal of the Optical Society of America A* 14, 1 (1997), 246–254. 2
- [LPT05] LI H., PELLACINI F., TORRANCE K. E.: A hybrid Monte Carlo method for accurate and efficient subsurface scattering. *Rendering Techniques (Proc. Eurographics Symposium on Rendering)* (2005), 283–290. 3
- [NNDJ12] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 4 (2012), 60:1–60:11. 2, 5, 8
- [PBPP11] PAJOT A., BARTHE L., PAULIN M., POULIN P.: Representativity for robust and adaptive multiple importance sampling. *IEEE Transactions on Visualization and Computer Graphics* 17, 8 (2011), 1108–1121. 6
- [PH10] PHARR M., HUMPHREYS G.: *Physically Based Rendering: From Theory To Implementation*, 2nd ed. Morgan Kaufmann, 2010. 1, 2
- [Spa66] SPANIER J.: Two pairs of families of estimators for transport problems. *SIAM Journal on Applied Mathematics* 14, 4 (1966), 702–713. 2
- [Sta95] STAM J.: Multiple scattering as a diffusion process. *Rendering Techniques (Proc. Eurographics Workshop on Rendering)* (1995), 41–50. 2
- [VG95] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for Monte Carlo rendering. *Computer Graphics (Proc. SIGGRAPH)*, 29 (1995), 419–428. 6
- [WJZ95] WANG L., JACQUES S., ZHENG L.: MCML: Monte Carlo modeling of light transport in multi-layered tissues. *Computer Methods in Programs and Biomedicine*, 8 (1995), 313–371. 6, 8
- [WZT\*08] WANG J., ZHAO S., TONG X., LIN S., LIN Z., DONG Y., GUO B., SHUM H.-Y.: Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1 (2008), 1–18. 2
- [YZXW12] YAN L.-Q., ZHOU Y., XU K., WANG R.: Accurate translucent material rendering under spherical Gaussian lights. *Computer Graphics Forum (Proc. Pacific Graphics)* 31, 7 (2012), 2267–2276. 2, 7