

An Empirical BSSRDF Model

Craig Donner* Jason Lawrence† Ravi Ramamoorthi ‡
 Toshiya Hachisuka§ Henrik Wann Jensen§ Shree Naray*
 *Columbia University †University of Virginia ‡UC Berkeley §UC San Diego

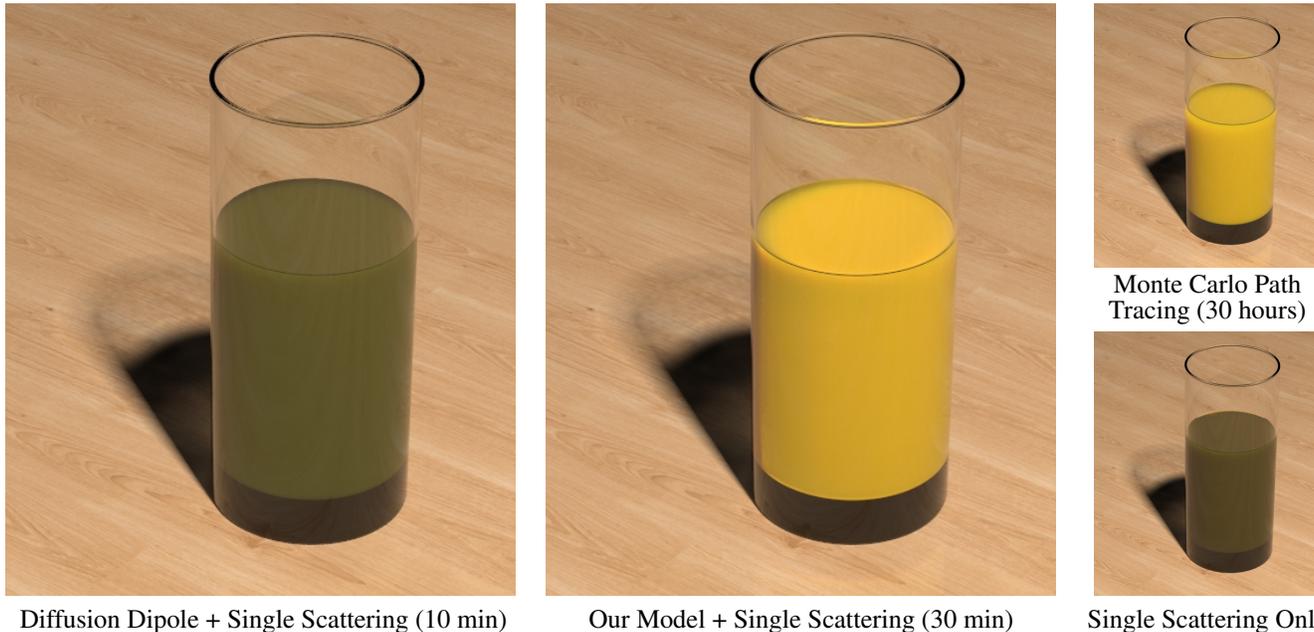


Figure 1: Although the appearance of orange juice is dominated by low-order scattering events, it is not accurately predicted by a single scattering model alone (lower right). Adding the contribution from high-order multiple scattering using the diffusion dipole (left) still fails to capture these effects and produces visible color artifacts. Numerical methods such as Monte Carlo path tracing (upper right) or photon mapping are accurate, but do not provide an explicit model of the BSSRDF and require long rendering times. Our proposed model is compact, efficient to render and can accurately express the complex spatial- and directionally-dependent appearance of these types of materials.

Abstract

We present a new model of the homogeneous BSSRDF based on large-scale simulations. Our model captures the appearance of materials that are not accurately represented using existing single scattering models or multiple isotropic scattering models (e.g. the diffusion approximation). We use an analytic function to model the 2D hemispherical distribution of exitant light at a point on the surface, and a table of parameter values of this function computed at uniformly sampled locations over the remaining dimensions of the BSSRDF domain. This analytic function is expressed in elliptic coordinates and has six parameters which vary smoothly with surface position, incident angle, and the underlying optical properties of the material (albedo, mean free path length, phase function and the relative index of refraction). Our model agrees well with measured data, and is compact, requiring only 250MB to represent the full spatial- and angular-distribution of light across a wide spectrum of materials. In practice, rendering a single material requires only about 100KB to represent the BSSRDF.

1 Introduction

Light propagates into and scatters within all non-metallic materials. This *subsurface scattering* is common in many liquids—such as orange juice, coffee or milk, and in solids—such as gemstones, leaves, wax, plastics and skin. It gives materials their characteristic colors, and provides a soft, translucent appearance. Accurate and compact models of the way light interacts with these materials are necessary to efficiently render them.

Light scattering in translucent materials is described by the bidirectional scattering surface reflectance distribution function S (the BSSRDF [Nicodemus et al. 1977]). The BSSRDF defines the general transport of light between two points and directions as the ratio of the radiance $L_o(\vec{x}_o, \vec{\omega}_o)$ exiting at position \vec{x}_o in direction $\vec{\omega}_o$ to the radiant flux $\Phi_i(\vec{x}_i, \vec{\omega}_i)$ incident at \vec{x}_i from direction $\vec{\omega}_i$:

$$\frac{dL_o(\vec{x}_o, \vec{\omega}_o)}{d\Phi_i(\vec{x}_i, \vec{\omega}_i)} = S(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o | \sigma_s, \sigma_a, g, \eta), \quad (1)$$

where S depends on the optical properties of the material—the scattering and absorption coefficients σ_s and σ_a , the relative index of refraction η , and $g \in [-1 : 1]$ which parameterizes the anisotropy of the phase function.

1.1 Related Work

Numerical techniques such as Monte Carlo path tracing [Kajiya 1986; Jensen et al. 1999] are capable of simulating general BSSRDFs. However, these methods are expensive, often requiring hours to days of processing. Scattering equations may also be used in this context [Pharr and Hanrahan 2000], but are computationally

expensive to evaluate. Photon mapping [Jensen 1996] can render many BSSRDFs but becomes expensive in both time and space for highly scattering materials. Furthermore, these techniques do not explicitly model the BSSRDF. Rather, determining the fraction of light that is transported between any pair of points requires a complete simulation.

A related set of techniques focus on simulating participating media (see Cerezo et al. [2005] for a survey), though many have difficulty handling refraction at an interface. In particular, Premoze et al. [2003] use an approximate path integral formulation to identify the most probable paths of light through a medium to efficiently render a wide range of scattering materials. More similar to our work, Premoze et al. [2004] analyze light transport in materials based on Monte Carlo simulations. They use the path integral technique to compute the contribution of collimated beams by first precomputing the reduced intensity within the scattering volume, and blurring this response using gaussian point spread functions to approximate the spatial and angular spread of light. As before, these methods do not give an explicit representation of the BSSRDF which is the focus of this paper.

Existing analytic models of the BSSRDF only apply to two classes of materials. When light scatters exactly once and at a single point, this *single scattering* has a closed-form analytic solution [Blinn 1982; Hanrahan and Krueger 1993]. Such models produce highly directional effects since the exitant light is assumed to be scattered directly from propagating beams of light. At the other end of the spectrum are highly scattering materials. The BSSRDF in these cases is often modeled as the superposition of a single scattering term and a diffuse term [Jensen et al. 2001]. This *diffusion approximation* [Stam 1995] is common and can be applied on its own to materials that have negligible low-order scattering [Jensen and Buhler 2002; Donner and Jensen 2005]. However, this approximation assumes light is scattered isotropically and produces incorrect predictions when low-order scattering is significant, such as in materials like orange juice (see Figure 1). These materials exhibit significant absorption of the light as it propagates through the material, so much of the energy is scattered back into the environment near the point of incidence. This light exits in areas and directions outside of the single scattering regime, but not in the high-order multiple scattering regime. As a result, the angular distribution of exitant light is not predicted well by single scattering, diffusion, or their sum.

Hybrid methods attempt to combine the simplicity of diffusion with the accuracy of general numerical techniques. Donner and Jensen [2007] introduce a method to model asymmetric diffuse reflectance by sampling beams of light using diffusion sources seeded by single-step photon tracing. This method assumes that the material is highly scattering and neglects near-source and directional effects. Li et al. [2005] couple path tracing with the diffusion approximation for long path lengths, but in materials with moderate scattering or non-trivial absorption the path tracing step becomes computationally intensive. Furthermore, since these methods rely on numerical simulation, they too fail to provide an explicit model of the BSSRDF.

Measuring the appearance of translucent materials is also a difficult task. Goesele et al. [2004], Tong et al. [2005], and Peers et al. [2006] measure point-to-point transport, but ignore the angular distribution of the exitant light. Also, their data is only suitable for reproducing the particular materials measured. Narasimhan et al. [2006] use dilution to measure the optical properties of a variety of scattering materials, but rely on general numerical techniques for rendering. Instead of measuring a range of materials, we opt to simulate them instead and derive an analytic model based on these simulations. In this regard, our approach is similar to the virtual

gonioreflectometry technique of Westin et al. [1992] for analyzing the reflectance of complex opaque surfaces.

Bouthors et al. [2008] adopted a similar empirical approach to ours to study the light transport within slabs of clouds using simulations that assume a fixed albedo and phase function. They propose a set of analytic functions to model the aggregate reflectance observed in these simulations, which allows efficiently rendering clouds with arbitrary shapes. Although we also use a large-scale Monte Carlo simulation to study the internal scattering of materials, we propose an explicit representation of the full BSSRDF that characterizes the transport between arbitrary pairs of surface locations over a large range of materials, and considers refraction at the material boundary.

1.2 Our Approach

We propose an analytic model of the BSSRDF that applies to a wide range of materials. Our model is phenomenological and derived from a large-scale simulation of the subsurface light transport for a range of optical properties and geometric configurations. Although our approach was inspired by data-driven reflectance models [Ward 1992; Dana et al. 1999; Matusik et al. 2003], our choice to simulate these effects avoids a difficult acquisition task and provides greater control over the materials we consider. Further, we validate our simulation and final model using measured data.

Although the full BSSRDF in Equation 1 is a 12D function (4 spatial parameters, 4 angular parameters and 4 optical parameters), we make the common assumption of a spatially uniform, homogeneous semi-infinite material. This reduces the dimensionality of the BSSRDF to 8D (discussed in the next section) and makes exploring this space more feasible.¹ We use an efficient photon tracing technique and a cluster of computers to reconstruct the 2D hemispherical distribution of exitant light over a dense sampling of the remaining geometric and optical variables (a 6D space). This dataset captures the complete spatial and angular appearance of a wide range of translucent materials. Although this required many months of processing, it only needs to be created once and is an important contribution of this work.

Based on an analysis of this simulated data, we propose an analytic function expressed in elliptic coordinates with six fit parameters that accurately captures the features of these hemispherical distribution functions. We demonstrate that this function fits the simulated data well and, in turn, agrees with measured data, including highly scattering materials such as milk and wax which exhibit clear non-diffuse and anisotropic behaviors. Importantly, the parameters of this function vary smoothly over the 6D space of remaining geometric and optical parameters. This allows tabulating and interpolating them away from the samples we considered in our simulation to provide a continuous representation of the full BSSRDF. Our final model consists of a table of parameter settings of this elliptic function (approx. 250MB of space, though in practice only a fraction of the data is needed at one time) that can be used directly for rendering. We present images rendered using this model that show complex directional effects such as glows around beams that would be impossible to render with existing diffusion-based techniques.

2 Simulating the Space of BSSRDFs

We used a Monte Carlo photon tracing algorithm to reconstruct slices of the BSSRDF. Similar to Hanrahan and Krueger [1993], we trace photons into a semi-infinite slab and record how much energy

¹ As we assume the material is homogeneous, for any fixed set of optical parameters, the function that is evaluated during rendering is only 3D.

they deposit at the surface. We discretize the hemisphere of outgoing light based on the exit trajectory of the photons. This requires significantly more photons and storage than a standard diffuse trace.

Assumptions and Parameterization: Recall that we assume a homogeneous, semi-infinite material which allows reducing the 12D BSSRDF in Equation 1 to an 8D function. Table 1 summarizes our notation and the geometry of our setup is illustrated in Figure 2.

Since the scattering and absorption coefficients σ_s and σ_a have infinite range, we choose to parameterize the BSSRDF in terms of the albedo $\alpha = \sigma_s / (\sigma_s + \sigma_a)$. Note that $\alpha \in [0..1]$. All materials with the same albedo have the same exitant response up to a distance scale which is determined by the mean free path length $\ell = 1 / (\sigma_s + \sigma_a)$. For example, two materials with optical coefficients (σ_s, σ_a) and $(2\sigma_s, 2\sigma_a)$ differ only in terms of ℓ and $\ell/2$. Although technically there is a chance a photon may exit the material at any distance from the incident beam, overall exitant intensity falls-off rapidly with increasing distance from the source. Therefore, we can consider only the local region around \vec{x}_i without missing a significant portion of the BSSRDF. These choices allow replacing the optical properties $(\sigma_s, \sigma_a, g, \eta)$ with (α, g, η) and provide a finite usable range for all parameters that is suitable for sampling.

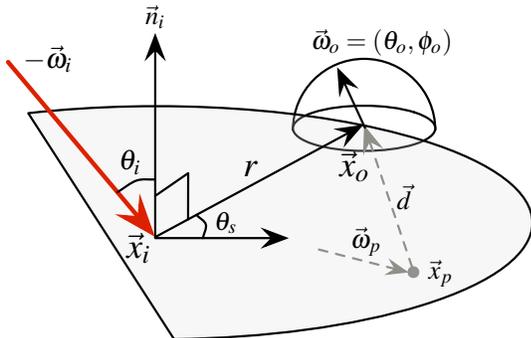
Just as with isotropic BRDFs, we assume that the angular dependence of the BSSRDF depends on only 3 variables $(\theta_i, \theta_o, \phi_o)$. Similarly, for a homogeneous BSSRDF the amount of light exchanged between two surface points \vec{x}_i and \vec{x}_o depends only on their relative positions. This allows parameterization of the spatial dimensions using 2D polar coordinates with the angle θ_s and radial distance $r = \|\vec{x}_o - \vec{x}_i\|$. Note that r is measured in mean free path lengths. Under the above assumptions, we simplify the BSSRDF as:

$$S(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o | \sigma_s, \sigma_a, g, \eta) \approx S(\theta_i, r, \theta_s, \theta_o, \phi_o | \alpha, g, \eta). \quad (2)$$

We next describe a photon tracing algorithm for reconstructing slices of this form of the BSSRDF. In Section 4 we introduce our new BSSRDF model.

2.1 Simulation Method

Our photon tracing algorithm relies on sampling the probable paths of light within a material. We discretize the exitant direction of each photon at the surface and accumulate its power into bins. However, to accurately resolve the complete hemispherical distribution of outgoing light a large number ($\sim 10^{12}$) of photons are required, particularly in the case of strong anisotropic scattering (i.e., $|g| > 0.5$).



$$S(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o | \sigma_s, \sigma_a, g, \eta) \approx S(\theta_i, r, \theta_s, \theta_o, \phi_o | \alpha, g, \eta).$$

Figure 2: Geometric setup for our simulation and model. A collimated beam of light arriving at point \vec{x}_i from direction $\vec{\omega}_i$ (shown in red) makes angle θ_i with the surface normal. A photon from this beam that arrives at point \vec{x}_p propagating in direction $\vec{\omega}_p$ has an exitant trajectory \vec{d} towards the point $x_o = (r, \theta_s)$ on the surface. This particular photon path exits the material in direction $\vec{\omega}_o = (\theta_o, \phi_o)$.

S	BSSRDF
\vec{x}_i	Surface location of incident light (source)
\vec{n}_i	Surface normal at \vec{x}_i
$\vec{\omega}_i$	Incident direction
θ_i, ϕ_i	Incident elevation and azimuthal angles
Φ_i	Incident radiant flux at \vec{x}_i
\vec{x}_o	Surface location of exitant light
\vec{n}_o	Surface normal at \vec{x}_o
r, θ_s	Polar coordinates of x_o w.r.t. source
$\vec{\omega}_o$	Exitant direction
θ_o, ϕ_o	Exitant elevation and azimuthal angles
\vec{x}_p	Location of a scattering event inside material
$\vec{\omega}_p$	Direction of incoming photon at x_p
Φ_p	Power of photon arriving at \vec{x}_p
Φ_e	Power of photon leaving \vec{x}_p arriving at \vec{x}_o
\vec{d}	Unit vector from \vec{x}_p to \vec{x}_o
\hat{s}	Single scattering plane
σ_s	Scattering coefficient
σ_a	Absorption coefficient
α	Albedo
p	Phase function
g	Anisotropy factor in phase function
ℓ	Mean free path length
η	Relative index of refraction
F_i	Fresnel transmittance
x_d, y_d	Projection of $\vec{\omega}_o$ onto unit disc
v, μ	Elliptic coordinates of (x_d, y_d)
$\vec{\omega}_{\text{peak}}$	Direction of maximum exitant intensity
\hat{s}_p	Plane of symmetry of exitant distribution
H	Proposed analytic distribution function
$\mathbf{a}^\pm, \mathbf{b}^\pm$	Focal points in elliptic coordinate system
$\mathbf{k}_s, \mathbf{k}_e, \mathbf{k}_c$	Parameters in analytic distribution function

Table 1: Notation used in this paper.

We emit photons along a collimated beam incident on the slab from direction $-\vec{\omega}_i$, such that it makes angle θ_i with the surface normal \vec{n} (see Figure 2). These photons refract into the material and propagate a distance d before being scattered or absorbed:

$$d = \frac{-\log \xi}{\sigma_s + \sigma_a}, \quad (3)$$

where $\xi \in [0..1]$ is a uniformly distributed random number.

Reusing Photon Paths: To reduce the number of photons traced, and thus increase the speed of our simulation, we compute the power Φ_e that a photon propagating in direction $\vec{\omega}_p$ incident on position \vec{x}_p would contribute at the surface point \vec{x}_o , if it were scattered directly there (see Figure 2):

$$\Phi_e = \alpha p(\vec{d} \cdot \vec{\omega}_p; g) e^{-(\sigma_s + \sigma_a) \|\vec{x}_o - \vec{x}_p\|} F_i(\vec{\omega}_o; \eta) \Phi_p, \quad (4)$$

where Φ_p is the power of the photon at \vec{x}_p , $\vec{d} = \frac{\vec{x}_o - \vec{x}_p}{\|\vec{x}_o - \vec{x}_p\|}$ is the normalized vector from \vec{x}_p to \vec{x}_o , p is the phase function, and F_i is the Fresnel transmittance at the surface. If internal reflection occurs,² then the photon makes no contribution. Otherwise, Φ_e is added to the exitant radiance at \vec{x}_o . We compute the contribution of each photon path to every exitant point we consider during our simulation. This results in a consistent estimate of the exitant radiance. To normalize the final radiance values, the power of each photon is scaled by the inverse of the number of photons traced.

² Internal reflection occurs when the exitant trajectory of the photon (relative to the normal) is greater than the critical angle: $\vec{d} \cdot \vec{n}_o > \sin^{-1} \eta$, where \vec{n}_o is the surface normal at \vec{x}_o and η is the ratio of the indices of refraction. Note that $\sin^{-1} \eta$ is the critical angle, and $\eta < 1$ from the point of view of the photon.

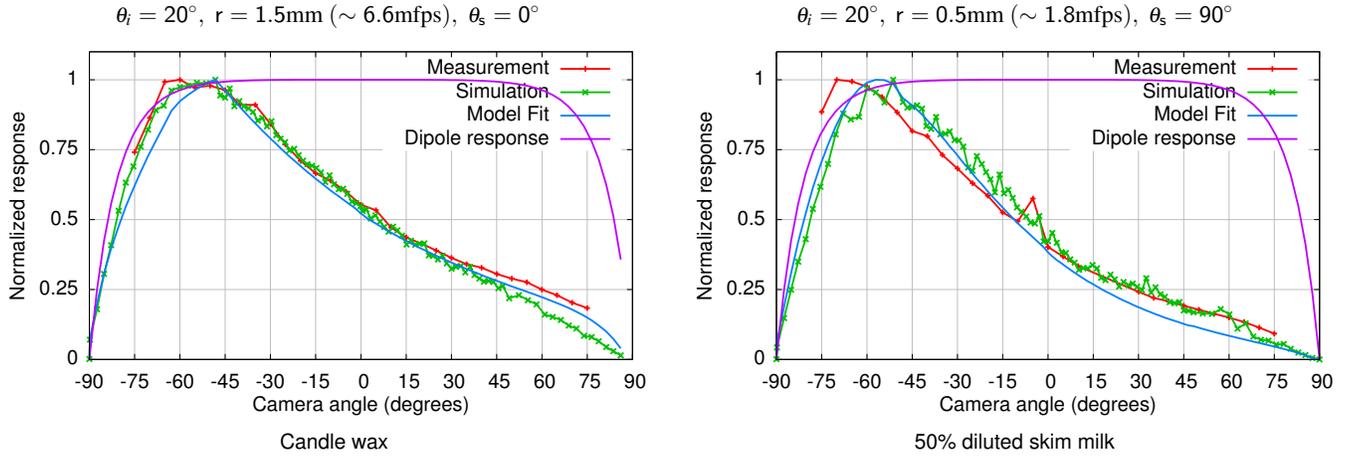


Figure 3: Slices of the 2D hemisphere of exitant light at different surface positions and different angles of incidence for wax (left) and 50% diluted skim milk (right). A diffusion approximation predicts a flat, diffuse response over the camera angle, whereas our simulated data and model agree well with measurements. We approximated the optical properties of milk as $\sigma_s = 1.165$, $\sigma_a = 0.0007$ ($\alpha \sim 0.99$, $\ell = 0.86\text{mm}$), $g = 0.7$, $\eta = 1.35$ [Joshi et al. 2006], and of wax as $\sigma_s = 1$, $\sigma_a = 0.5$ ($\alpha = 0.67$, $\ell = 0.67\text{mm}$), $g = 0$, $\eta = 1.4$.

Variance Reduction: We use Russian Roulette to determine whether this photon, having just traveled the distance d , is scattered or absorbed. If it is scattered, then we compute its new trajectory by importance sampling the Henyey-Greenstein phase function [Jensen 2001]. Therefore, we do not modify the weight or power of a photon unless it probabilistically reaches the surface, at which point it is internally reflected after having applied the Fresnel term. We continue to trace each photon until it is absorbed.

2.2 A Database of BSSRDFs

Our photon tracing algorithm allows simulating the full 8D BSSRDF in Equation 2 to construct a database we will use to develop and validate our model. To accurately resolve near-source and directional effects, we densely sample the 2D exitant hemisphere (θ_o, ϕ_o) at ~ 6200 uniformly distributed directions at each exitant location (r, θ_s) on the surface. Since the response of a material type does not depend on the mean free path length ℓ , we fix $\ell = 1\text{mm}$ in all of our simulations. During rendering, we account for the true mean free path length of a specific material by simply applying the appropriate scale factor (Section 5). We sample the remaining BSSRDF dimensions as follows:

$$\begin{aligned} \theta_i &\in \{0, 15, 30, 45, 60, 70, 80, 88\} \\ r &\in \{.01, .05, .1, .2, .4, .6, .8, 1, 2, 4, 8, 10\} \\ \theta_s &\in \{0, 15, 30, 45, 60, 75, 90, 105, 120, 135, 150, 165, 180\} \\ \alpha &\in \{0.01, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99\} \\ g &\in \{-0.9, -0.7, -0.5, -0.3, 0, 0.3, 0.5, 0.7, 0.9, 0.95, 0.99\} \\ \eta &\in \{1, 1.1, 1.2, 1.3, 1.4\} \end{aligned}$$

Note that r , the distance in mean free paths from the incident beam, increases at an exponential rate to account for the fact that the total exitant intensity diminishes along a similar trend. The other patterns provide a mostly uniform sampling of the remaining dimensions. Note that we sample this space more finely near extremely low and high albedo and for very high forward scattering.

These sampling patterns produce about 750,000 exitant hemispheres, for a total of $\sim 4.7 \times 10^9$ points in this 8D space. Generating this amount of data poses a significant challenge since obtaining acceptable noise levels for just one hemispherical slice can require many hours of processing time. We performed these simulations on 15 Dual-Socket Quad-Core 2.33GHz Intel[®] Xeon[®] machines, for a total of 120 processing cores. Reconstructing the exitant hemi-

spherical distribution for a single set of optical properties requires about 60MB of storage—the entire database is roughly 36GB.

2.3 Validation

To verify the accuracy of our simulation and the model we propose in Section 3, we measured the distribution of exitant light for two materials: diluted milk and wax. These materials were illuminated using a red (635nm) laser dot of diameter $\sim 1\text{mm}$ at a fixed angle of $\theta_i = 20^\circ$ and were imaged using a monochromatic QImaging Retiga 4000R camera with a 60mm lens attached to a motorized gantry with an angular precision of $\sim 0.01^\circ$. We moved the camera along an arc at a fixed stand-off distance of 1.5m from the sample and recorded a high-dynamic range image every 5° over the range $\theta_o \in [-75^\circ, 75^\circ]$. This set of images provides 1D angular slices of the exitant hemispherical distributions for many points on the material surface. We repeated this procedure for two arcs: one parallel to the plane of incident light and one perpendicular to this plane. The camera was photometrically calibrated and the resulting images were aligned using standard chart-based calibration techniques [Zhang 1999].

Figure 3 compares measured slices of the exitant light to predictions made by our simulation (Section 2.1), a standard diffusion approximation [Jensen et al. 2001] and our proposed model (we discuss the “Model Fit” curves in Section 4.1). Note that due to the angles involved in the measurement, the plane of single scattering was explicitly excluded. The optical properties we used in these simulations are reported in the caption. Note that because the diffusion approximation predicts a diffuse response, modulated only by Fresnel transmittance, it does not capture the clear asymmetry in these distributions. Our simulation closely follows the measured data. At these distances from the source there is some visible noise in the simulations as the paths are relatively long for high-albedo materials. Unfortunately, since the size of the laser dot is large relative to the mean free path of these materials, it was difficult to obtain accurate measurements very close to the source. Although these plots show only two configurations of r and θ_s , they are representative of the close agreement we observed over the entire material surface.

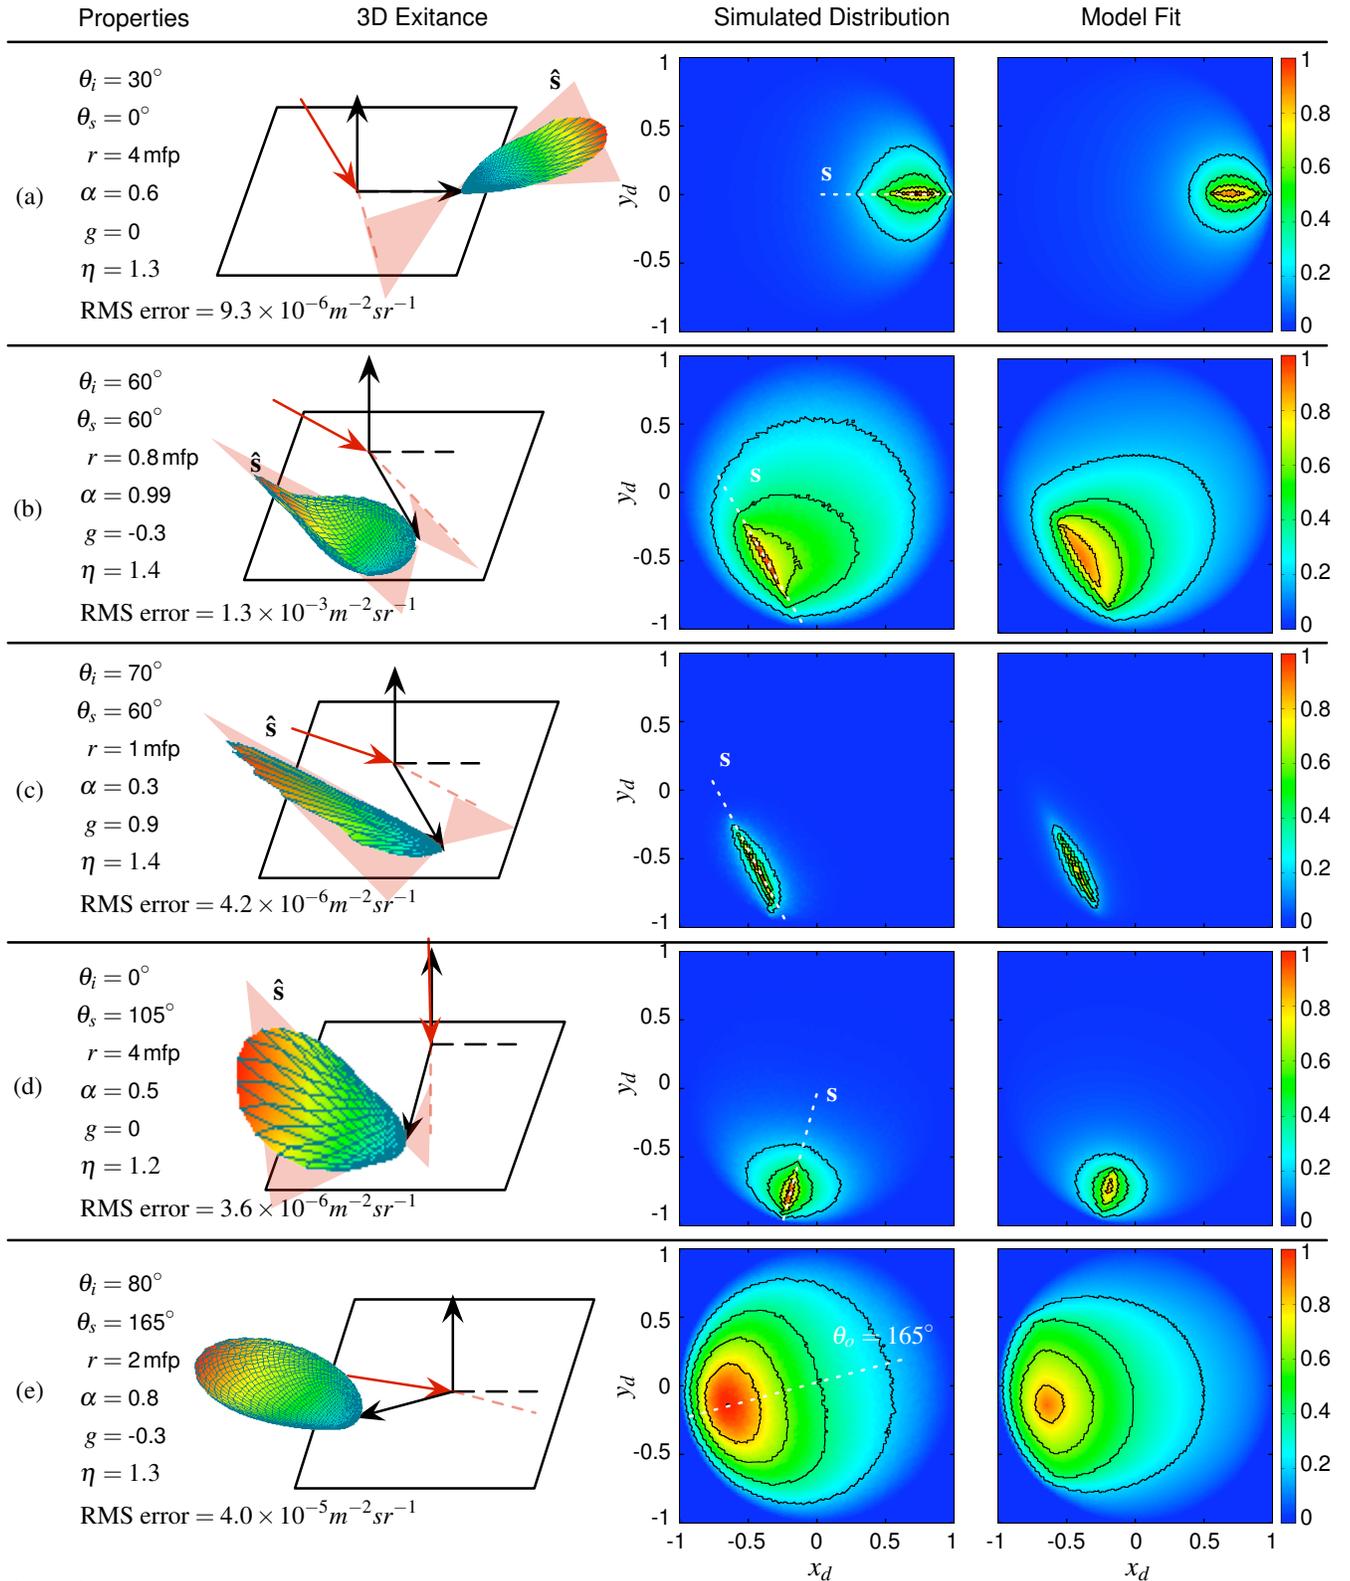


Figure 4: Representative 2D slices of $S(\theta_i, r, \theta_s, \theta_o, \phi_o | \alpha, g, \eta)$ defined over θ_o and ϕ_o . The left column lists the optical properties (α, g, η), the geometric setup (θ_i, r, θ_s), and the RMS error of the fit. The second column contains diagrams that show the incident beam in solid red, the propagating beam inside the material in dashed red, and the simulated distribution of the exitant radiance. The last two columns show false-color plots of these hemispherical distributions projected onto the disc according to Equation 5. The “Simulated Distribution” plots were produced with our photon tracing algorithm and the “Model Fit” plots correspond to our proposed analytic model. To aid comparison, all plots are normalized to unity. The single scattering plane \hat{s} is shown in the cases where the lobe is symmetric about this plane. Note that all of these distributions, except for (e), exhibit this type of symmetry while (e) is symmetric about the plane formed by $\vec{x}_i - \vec{x}_o$ and the surface normal \vec{n} .

3 Data Analysis

The simulation method described previously produces 2D slices of the BSSRDF $S(\theta_i, r, \theta_s, \theta_o, \phi_o | \alpha, g, \eta)$ over outgoing angles θ_o and ϕ_o . In this section we describe a number of key observations about these functions that will serve as the basis of our model presented in Section 4. Specifically, we are interested in understanding how these hemispherical functions vary with respect to the optical properties α , g , and η , along with their dependence on the incident angle of the beam θ_i and surface position (r, θ_s) .

Figure 4 visualizes several hemispherical slices produced by our simulation at different parameter settings. The leftmost column lists the optical properties, geometric configuration, and RMS error of our fits (described in the next section). Next to these values is a 3D visualization of the simulated exitant light, along with the incident beam (in red) and the path it takes into the material after being refracted at the surface (dashed red). The single scattering plane \hat{s} is also shown. This plane traces an arc across the exitant hemisphere, and is shown as a wedge emanating from the refracted beam. Because the overall intensity of outgoing light is proportional to the distance from the surface location \vec{x}_o these plots are normalized for visualization purposes. These types of visualizations are helpful in understanding how the shapes of these distributions are related to the angular configuration of the setup, and reveal important symmetries in the data.

The third column in Figure 4 shows projections of these hemispherical functions onto the unit disc. The color scale, shown at right, assigns higher intensity values red and lower values blue—this same scheme is used in the 3D plots as well. We map the hemisphere to the unit disc using a standard projection:

$$x_d = \frac{2\theta_o}{\pi} \cos \phi_o \quad \text{and} \quad y_d = \frac{2\theta_o}{\pi} \sin \phi_o. \quad (5)$$

The variety of shapes present in these 2D distributions indicates a complex relationship with the remaining parameters of the BSSRDF. Furthermore, the presence of high-frequency features in these distributions means that they would not be suitable for rendering directly. Linearly combining these slices in order to interpolate regions of the BSSRDF domain that were not directly simulated would produce significant errors. This motivates our goal of providing a compact analytic function that captures the shape of these distributions. Interpolating the parameters of this analytic function provides a reliable and efficient way of modeling the continuous 8D BSSRDF domain.

3.1 Phenomenological Observations

Based on this collection of simulated data we make the following key observations:

Anisotropy: The exitant distribution of light is not diffuse. We observe that even highly scattering materials, such as in Figure 4b, exhibit a non-diffuse and anisotropic shape. This follows from the fact that as light is scattered away from the propagating beam of light it is focused near the single scattering plane. Therefore, the majority of the exitant energy is due to low-order scattering events that occur before photons have lost much of their directionality.

Peak Direction and Kurtosis: With a few exceptions discussed in Section 7, these distributions almost always contain a single peaked lobe. We also observe that the direction of this peak depends on θ_i . We attribute this to the fact that at steeper angles more photons propagate a further distance across the surface than below the surface before being scattered. The shape of this lobe also depends on the surface position (r, θ_s) since points near \hat{s} will receive a higher contribution of light. Additionally, the optical properties of the material clearly affect the shape of these exitant lobes. Higher

values of α produce more scattering and thus a wider lobe. With larger magnitudes of g , however, light is less likely to scatter outside the direction of the propagating beam and thus produces a tighter, sharper peak. This was especially true for forward scattering materials such as Figure 4c. Finally, since the direction of propagation of the beam depends on η , it also affects the orientation of the lobe.

Lobe Asymmetry: These lobes are typically not rotationally symmetric about their axes and the degree of asymmetry depends primarily on the angle θ_s . As photons propagate into the material near the refracted beam and then scatter away towards the surface, points on the surface receive the strongest contribution from regions near the propagating beam. The distributions at exitant points further away from the propagating beam exhibit greater asymmetries since the photons contributing to these locations have traveled a greater distance. This asymmetry is particularly visible in Figure 4b.

Lobe Shape: The lobe is often aligned with the single scattering plane \hat{s} , as in Figure 4a-d, indicating a strong contribution from low-order scattering. However, when the single scattering trajectory exceeds the critical angle (i.e. when $\vec{d} \cdot \vec{n} > \sin^{-1} \eta$), total internal reflection occurs and produces lobes with a wider, less peaked shape, but ones that are still not entirely diffuse. This is visible in Figure 4e. Here, the exitant lobe tends to be symmetric about the plane aligned with θ_s , as shown.

Elliptical to Circular Isocontours: The black isocontour lines shown in Figure 4 reveal that these lobes have an elliptical shape near the peak which transitions to a more circular shape further away from the peak. We attribute this to the number of times light was scattered by the material before arriving at these different exitant directions. Light that exits near \hat{s} has likely only scattered two or three times giving the exitant distribution a sharp elliptical contour along the projection of the propagating beam onto the hemisphere. Photons that travel longer paths contribute less power and are likely to exit in more uniform directions, resulting in a more diffuse distribution with circular contours.

4 An Empirical BSSRDF Model

Our goal is to model the 2D distribution of exitant light of the BSSRDF over the possible range of geometric and optical parameters. The observations made in the previous section indicate that functions traditionally used to model BRDFs, such as cosine or Gaussian lobes, are not applicable in this case. Because these lobes also tend to be sharply peaked, generic basis functions such as spherical or zonal harmonics are also unsuitable. Note that in this section, parameters to our empirical function are in bold; vectors are indicated with hats or arrows.

Indeed, the shape, position and size of the dominant lobe of exitant light has a complex relationship with the underlying optical and geometric properties. An important observation, however, is that cross-sectional contours of these lobes change from elliptical to circular away from the peak direction. This motivates our choice to define a function for these lobes in *elliptic coordinates* (μ, ν) [Korn and Korn 2000]. This coordinate system defines a set of confocal ellipses that have high eccentricity near the origin and slowly become more circular further away (see Figure 5).

The coordinates μ and ν are related to (x_d, y_d) , the projection of $\vec{\omega}_o$ onto the unit disc given in Equation 5, according to:

$$x_d = \begin{cases} \mathbf{a}^+ \cosh \mu \cos \nu, & x_d \geq 0 \\ \mathbf{a}^- \cosh \mu \cos \nu, & x_d < 0 \end{cases}, \quad y_d = \begin{cases} \mathbf{b}^+ \sinh \mu \sin \nu, & y_d \geq 0 \\ \mathbf{b}^- \sinh \mu \sin \nu, & y_d < 0 \end{cases} \quad (6)$$

and

$$\frac{x_d}{\mathbf{a}^\pm} + i \frac{y_d}{\mathbf{b}^\pm} = \cosh(\mu + i \nu). \quad (7)$$

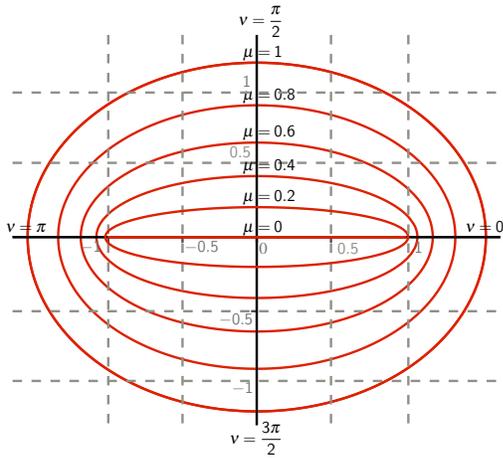


Figure 5: The elliptic coordinates describe a set of confocal ellipses. They provide a convenient system for representing our empirical BSSRDF model.

Although \mathbf{a}^+ , \mathbf{a}^- , \mathbf{b}^+ and \mathbf{b}^- are typically equal, allowing them to vary independently provides greater control over the shape of a lobe defined with respect to these coordinates. Note that when $\mu = 0$, \mathbf{a}^+ and \mathbf{a}^- define the distances of the elliptical focal points from the origin.

Isocontours and Alignment: Elliptic coordinates are well suited for capturing the “elliptical-to-circular” trend we observed in our simulations. Considering the projection of these slices onto the disc, we place the origin of this coordinate system at the peak direction $\vec{\omega}_{\text{peak}}$ (location of maximum exitant intensity) and define the x -axis to lie along the plane of symmetry $\hat{\mathbf{s}}_p$. This plane is either the single scattering plane $\hat{\mathbf{s}}$, or the plane defined by $\vec{x}_i - \vec{x}_o$ and the surface normal \vec{n} for distributions dominated by internal reflection as previously discussed. We propose the following analytic function of these 2D exitant distributions:

$$H(\vec{\omega}_o; \vec{\omega}_{\text{peak}}, \hat{\mathbf{s}}_p, \mathbf{k}_s | \Gamma) = \mathbf{k}_s e^{-\mathbf{k}_e \mu - \mathbf{k}_c \chi} F_t(\vec{\omega}_o, \eta), \quad (8)$$

where F_t is the Fresnel transmittance in the outgoing direction $\vec{\omega}_o$, $\chi = \sqrt{x_d^2 + y_d^2}$ is the distance on the unit disc from the origin $\vec{\omega}_{\text{peak}}$, \mathbf{k}_s is the overall scale of the distribution, defined such that the fit intensity matches the data at $\vec{\omega}_{\text{peak}}$, and $\Gamma = \{\mathbf{a}^+, \mathbf{a}^-, \mathbf{b}^+, \mathbf{b}^-, \mathbf{k}_e, \mathbf{k}_c\}$ is a 6D vector that collects the free parameters in this function. These control the lobe’s shape, degree of asymmetry, anisotropy and kurtosis as explained below.

Asymmetry and Anisotropy: The parameters \mathbf{a}^+ , \mathbf{a}^- , \mathbf{b}^+ , and \mathbf{b}^- determine the asymmetry of the elliptical contours. When they are small, the lobe is focused; as they increase, the lobe becomes more diffuse. When they are unequal, the lobe is asymmetric and anisotropic.

Lobe Shape: The parameters \mathbf{k}_e and \mathbf{k}_c determine the shape of the cross-section of the lobe (i.e., capturing the transition from elliptical to circular). The parameter \mathbf{k}_e controls the falloff along the plane of symmetry $\hat{\mathbf{s}}_p$. Distributions with long, high peaks have lower \mathbf{k}_e , while tightly focused peaks require larger values. The radially symmetric shape is determined by \mathbf{k}_c . Although these parameters are correlated to one another, especially further from the peak direction, maintaining these degrees of freedom provides finer control over the precise shape of the lobe near $\vec{\omega}_{\text{peak}}$, where most of the exitant light is focused.

4.1 Fitting the Model to Simulated Data

We fit the six parameters in Γ to 2D slices produced by our simulation for each configuration of the optical and geometry properties using direct analysis. We found that general non-linear optimization routines were unnecessarily complex and often produced noisy fits which undermines our goal of smoothly interpolating these values over the full BSSRDF domain.

Our fitting procedure is iterative. We first record values of the exitant distribution along a set of isocontour levels, e.g. $c_1 = 10\%$, $c_2 = 20\%$, ..., $c_8 = 80\%$, and $c_9 = 90\%$ of the peak value. For simplicity, we only examine points along the x_d -axis as this simplifies the relationship between cartesian and elliptic coordinates (since $v = 0$ when $y_d = 0$):

$$\mu = \cosh^{-1} \left(\frac{x_d}{\mathbf{a}^\pm} \right), \quad y_d = 0. \quad (9)$$

At each contour sample, we calculate μ using Equation 9. The values of x_d at these samples are readily obtained. More than two samples results in an overconstrained linear system computed by partially inverting Equation 8:

$$-\log \left(\frac{c_i H(\vec{\omega}_{\text{peak}})}{F_t(\vec{\omega}_{\text{peak}})} \right) = \mathbf{k}_e \cosh^{-1} \left(\frac{x_d^\pm}{\mathbf{a}^\pm} \right) + \mathbf{k}_c x_d^\pm \quad (10)$$

We solve this system to obtain values of \mathbf{k}_e and \mathbf{k}_c that best fit these contour samples in the least squares sense. We repeat this procedure along the positive and negative sides of the x_d -axis, and retain the smaller of the two values for both \mathbf{k}_e and \mathbf{k}_c .

Given these values for \mathbf{k}_e and \mathbf{k}_c it is straightforward to solve for \mathbf{a}^\pm and \mathbf{b}^\pm from Equation 10. For example, for x_d^\pm :

$$\mathbf{a}^\pm = \frac{x_d^\pm}{\cosh \left(\frac{x_d^\pm \mathbf{k}_c - \log \left(\frac{c_i H(\vec{\omega}_{\text{peak}})}{F_t(\vec{\omega}_o)} \right)}{\mathbf{k}_e} \right)}. \quad (11)$$

The expression for y_d^\pm is similar. We solve for these values at each contour level, and choose the set of \mathbf{a}^\pm and \mathbf{b}^\pm which minimizes the total L_2 error over the hemisphere with respect to the original exitant distribution. We then re-estimate \mathbf{k}_e and \mathbf{k}_c using these new values of the other parameters and repeat this process ten times or until we observe the error change by less than 1%.

4.2 Final Model

For each slice of the BSSRDF computed in our simulation we store the best fit parameters of our model H in a large table, a total of $\sim 250\text{MB}$.³ Because the two dimensions of $\vec{\omega}_o$ are the most densely sampled in our simulation this is a significant reduction, from ~ 6200 points on the hemisphere to the set of inputs to H (11 values). Finally, note that in practice only a fraction of this data is used at any one time during rendering.

4.3 Model Accuracy

Figure 3 compares the best fits of our model to simulated data, measured data and an approximation of these angular distributions obtained with a standard dipole diffusion model. These results show strong agreement between measured and simulated data and our proposed model.

The fourth column of Figure 4 compares fits of our model H to simulated data for several representative exitant distributions. Although our model does not provide an exact match it successfully

³ This data can be downloaded, along with the raw simulated data and the measurements used in Figure 3, as auxiliary material from the ACM Digital Library at <http://portal.acm.org>.

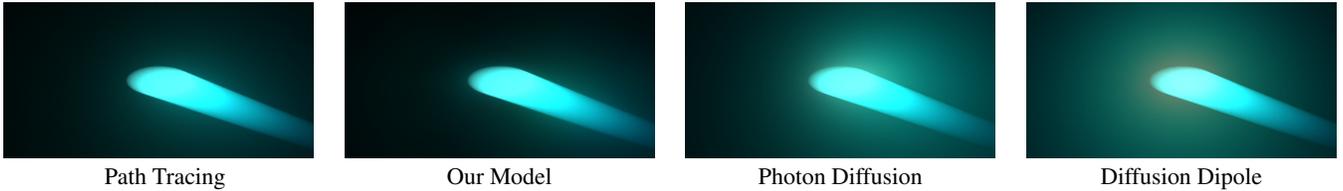


Figure 6: A beam of light striking the material surface at a 60° angle off the normal. The leftmost image was rendered using path tracing, and shows the expected glow around the propagating beam. The image rendered using our model closely agrees to the reference path-traced result. Both photon diffusion and the diffusion dipole incorrectly predict the angular distribution of energy within this material as well as the spectral distribution of emitted light.

captures the basic trends in these distributions over a wide range of optical properties and geometric configurations. In particular, note that H exhibits the characteristic “elliptical-to-circular” pattern.

We also computed the RMS error of our fits over the entire collection of simulated data. These errors have a mean of 3.2×10^{-3} and standard deviation of 9.0×10^{-4} . Note that these values have the same units as the BSSRDF of $[m^{-2}sr^{-1}]$. The RMS error of the fits shown in Figure 5 from (a) to (e) are 9.3×10^{-6} , 1.3×10^{-3} , 4.2×10^{-6} , 3.6×10^{-6} , and 4.0×10^{-5} , respectively. We also validated the process of interpolating the parameters of H to predict exitant distributions that were not directly simulated. We computed the RMS error between each simulated distribution and that produced by interpolating the parameters of our model using those values at its nearest neighbors. This is akin to a “leave one out” cross validation test. These errors had a mean of 4.9×10^{-4} with a standard deviation of 6.9×10^{-7} . We conclude that our model is capable of fitting individual slices with a high degree of numerical accuracy and that its parameters can be safely interpolated to reconstruct a continuous representation of the BSSRDF over the range of geometric and optical properties considered in our simulations.

As further validation of our model and to compare it to alternative methods, we produced renderings of a beam of light entering a material with $\alpha = (0.07, 0.53, 0.52)$ and $g = 0$ (see Figure 6). Although all four rendering methods capture single scattering well, both the diffusion dipole and photon diffusion inaccurately predict the color and intensity of multiply scattered light. Our method, on the other hand, captures the correct wavelength-dependence along with the directional effects of the internally scattered light.

5 Rendering with the Model

We sample the illumination incident on translucent materials using standard techniques (e.g. [Jensen et al. 2001]). To compute a set of surface locations with respect to a single shade point, we draw samples from a probability density function proportional to the exitant diffuse distribution based on the scale factor \mathbf{k}_s in H . Alternatively, we could use a hierarchical point-based approximation of the irradiance [Jensen and Buhler 2002] to perform this sampling. This would decrease rendering times without affecting the model itself.

Interpolation: We first reconstruct the BSSRDF that corresponds to the optical properties of the material we are interested in rendering. We linearly interpolate the parameters of H based on the closest neighboring points along the α , g , and η axes. This gives a compact (~ 100 KB) 3D reconstruction defined over θ_i , r , and θ_s .

For a shade point \vec{x}_o and lit point \vec{x}_i on the surface we compute r and θ_s with respect to an orthogonal coordinate system with x -axis along $\vec{\mathbf{X}} = \vec{\omega}_i - (\vec{\omega}_i \cdot \vec{n}_i)\vec{n}_i$. The surface normal at \vec{x}_i is \vec{n}_i which we assume points up and the z -axis of this coordinate system is defined

as $\vec{\mathbf{Z}} = \vec{\mathbf{X}} \times \vec{n}_i$. This construction leaves θ_s as the rotation:

$$\begin{aligned} \theta_i &= \cos^{-1}(\vec{\omega}_i \cdot \vec{n}_i) \\ r &= (\sigma_s + \sigma_a) \|\vec{x}_o - \vec{x}_i\| \\ \theta_s &= \tan^{-1} \left(\frac{(\vec{x}_o - \vec{x}_i) \cdot \vec{\mathbf{Z}}}{(\vec{x}_o - \vec{x}_i) \cdot \vec{\mathbf{X}}} \right). \end{aligned} \quad (12)$$

To determine the interpolated locations along the r dimension we compute the distance in terms of mean free paths by scaling by the inverse of the mfp length, i.e. $\sigma_s + \sigma_a$.

Handling Non-Planar Geometry: Once we have interpolated the parameters of our elliptic function as described above, we next determine the values of θ_o and ϕ_o to use when evaluating H . Because our model is only valid for semi-infinite materials, care must be taken when rendering arbitrary geometry (Section 7).

We define a separate coordinate system anchored at \vec{x}_o using the normal \vec{n}_o , $\vec{\mathbf{X}}_h = \vec{n}_o \times \vec{\mathbf{Z}}$ and $\vec{\mathbf{Z}}_h = \vec{\mathbf{X}}_h \times \vec{n}_o$, and derive θ_o , ϕ_o as:

$$\begin{aligned} \theta_o &= \cos^{-1}(\vec{\omega}_o \cdot \vec{n}_o) \\ \phi_o &= \tan^{-1} \left(\frac{(\vec{\omega}_o - (\vec{\omega}_o \cdot \vec{n}_o)\vec{n}_o) \cdot \vec{\mathbf{Z}}_h}{(\vec{\omega}_o - (\vec{\omega}_o \cdot \vec{n}_o)\vec{n}_o) \cdot \vec{\mathbf{X}}_h} \right). \end{aligned} \quad (13)$$

Using these values, we locate the nearest 8 points in the 3D dataset, and interpolate the values of the model parameters ($\vec{\omega}_{\text{peak}}, \hat{\mathbf{s}}_p, \mathbf{k}_s | \Gamma$). Finally, we calculate the value of the BSSRDF by evaluating H in Equation 8 using these interpolated values.

6 Results

All of the results in this paper were rendered on an Intel[®] Xeon[®] 2.33GHz Quad-Core processor and those produced with our model required less than one hour of processing time.

Figure 1 compares renderings of orange juice produced using our method to those produced with Monte Carlo path tracing and the diffusion dipole combined with a single scattering term [Jensen et al. 2001]. We used parameters $\sigma_s = (0.071, 0.1, 0.042)$, $\sigma_a = (0.093, 0.16, 1.15)$, $g = 0.9$, and $\eta = 1.3$, which produces a relatively low spectral albedo of $\alpha = (0.43, 0.38, 0.035)$. Since orange juice is highly forward scattering, simulating single scattering alone does not capture its true appearance. High-order multiple scattering is also not a dominant effect and the image rendered using the diffusion dipole is too dark. This is because the dipole model assumes that anisotropic scattering is balanced by many scattering events. However, the low albedo (i.e. high absorption) of this material causes most of the light to be absorbed before reaching this regime. Because our model (center) more accurately captures these low order scattering events, it matches the reference path traced image, but requires significantly less time to compute.

Figure 7 shows a green lozenge with significant translucency. This material has a spectral albedo of $\alpha = (0.16, 0.27, 0.15)$ and is forward scattering with $g = 0.5$. The scene has global illumination which can be seen in the bleeding of green light from the lozenge

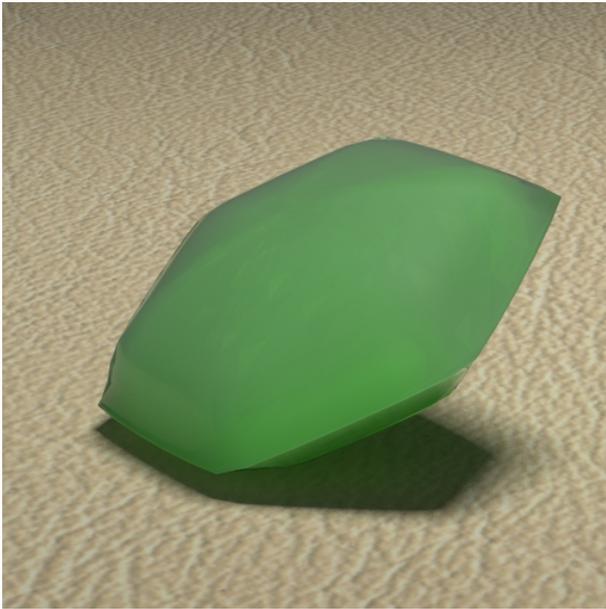


Figure 7: A translucent green lozenge with albedo $\alpha = (0.16, 0.27, 0.15)$. Note the global illumination between the subsurface scattering as predicted by our empirical BSSRDF model and the ground plane.

onto the ground plane. With our model, it is possible to efficiently render a much larger range of translucent materials than previously possible, even in scenes with global illumination.

To demonstrate the generality of our model we rendered the image shown in Figure 8 which combines several of the materials measured by Narasimhan et al. [2006]. These materials have significant absorption and anisotropic scattering which make them impossible to render correctly using standard diffusion-based techniques. The total amount of data required to render this image was about 300KB, as each material requires about 100KB for rendering.

7 Limitations

Our analytic formula of the exitant distribution of light H did fail to capture the proper response for some of the materials we observed. In particular, materials with strongly anisotropic backscattering ($g \ll 0$) can produce two exitant lobes (Figure 9). One is due to low-order scattering and is in the expected direction, but the other is caused by higher-order scattering. Clearly, as our model consists of only a single lobe it cannot handle these cases. However, because most real-world materials are forward scattering we concentrated on modeling the more prominent lobe and leave these cases for future study.

When light arrives at the surface along grazing angles internal refraction becomes more prevalent and our model is less accurate. Also, our assumption of a semi-infinite and flat surface overestimates the exitant light near corners and thin geometric features. This limitation is shared by previous work based on the diffusion dipole. Though in Section 5 we describe an approximate method for constructing a local coordinate frame based on the incident and exitant points on the surface, this method is not well-suited for optically thin materials, and does not account for the changes in internal scattering due to curved or cornered geometry. In the future we intend to investigate how best to overcome this limitation, such as using additional simulations similar to Bouthers et al. [2008]. The geometry has a significant influence on the appearance, and fur-

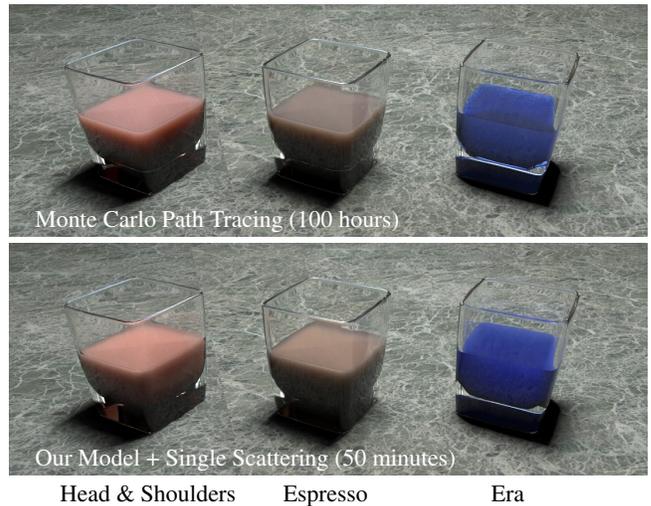


Figure 8: Renderings using parameters from [Narasimhan et al. 2006] to show that our model is capable of simulating a diverse range of materials.

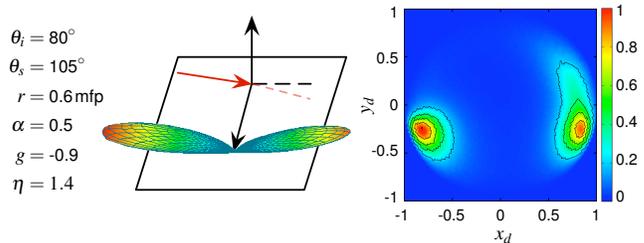


Figure 9: Our model is not able to represent highly backscattering materials that produce multiple lobes of scattered light.

ther investigation is needed to study precisely how local geometry affects the scattering of light, particularly in this regime of mid-albedo materials.

Further research is also warranted to determine a simpler relationship between the optical and geometric properties of the BSSRDF and the parameters in the analytic function H . While our own preliminary work in this direction has indicated that this mapping is complex and non-linear, developing an analytic relationship would shed further light on this important class of functions and increase the efficiency and accuracy of simulating these types of materials.

8 Conclusion

We presented an empirical model of the BSSRDF that is valid over a far wider range of angular configurations and material properties than existing analytic models. Our model captures both near-source and directional effects including the important contribution of low-order scattering. This model was derived from a large-scale simulation of the hemispherical distribution of light leaving a material's surface over a range of positions from the source, incident angles and underlying optical properties (scattering and absorption coefficients, phase function, and index of refraction). We presented an analytic function to approximate these hemispherical functions which is expressed in elliptic coordinates and has six parameters. We estimated the best-fitting parameters for our simulated data. Because these parameters vary smoothly with respect to the remaining degrees of freedom they may be interpolated between simulated locations to provide a compact yet continuous representation over the full BSSRDF domain. This allows generating realistic images of translucent materials with notoriously difficult optical proper-

ties. In particular, many of the results we reported would have been impossible to render using diffusion based methods and much less efficient with more general numerical integration techniques.

9 Acknowledgments

We thank the reviewers for their helpful input, and UCSD's FW-Grid project for access to their computational resources during the initial phases of this project. This work was supported in part by NSF grant 05-41259 on Fast and Accurate Volumetric Rendering of Scattering Phenomena in Computer Graphics, as well as NSF grants 03-25867, 04-46916, 07-01775, 07-01992, ONR Young Investigator Awards N00014-07-1-0900 and N00014-09-1-0741, and a Sloan Research Fellowship. Jason Lawrence acknowledges a NSF CAREER award 07-47220, NSF grant 08-11493 and an NVIDIA Professor Partnership Award.

References

- BLINN, J. F. 1982. Light reflection functions for simulation of clouds and dusty surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 82)*, vol. 16, 21–29.
- BOUTHORS, A., NEYRET, F., MAX, N., BRUNETON, E., AND CRASSIN, C. 2008. Interactive multiple anisotropic scattering in clouds. In *ISD '08: Proceedings of the 2008 symposium on Interactive 3D graphics and games*, 173–182.
- CEREZO, E., PEREZ-CAZORLA, F., PUEYO, X., SERON, F., AND SILLION, F. 2005. A survey on participating media rendering techniques. *The Visual Computer*.
- DANA, K., GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graphic.* 18, 1, 1–34.
- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Trans. Graphic.* 24, 3, 1032–1039.
- DONNER, C., AND JENSEN, H. W. 2007. Rendering translucent materials using photon diffusion. In *Rendering Techniques*, 243–251.
- GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND SIEDEL, H.-P. 2004. DISCO: acquisition of translucent objects. *ACM Trans. Graphic.* 23, 3, 835–844.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 1993*, 164–174.
- JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graphic.* 21, 576–581.
- JENSEN, H. W., LEGAKIS, J., AND DORSEY, J. 1999. Rendering of wet materials. In *Rendering Techniques*, 273–282.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, 511–518.
- JENSEN, H. W. 1996. Global illumination using photon maps. In *Rendering Techniques*, 21–30.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. AK Peters.
- JOSHI, N., DONNER, C., AND JENSEN, H. W. 2006. Noninvasive measurement of scattering anisotropy in turbid materials by nonnormal incident illumination. *Opt. Lett.* 31, 936–938.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proceedings of SIGGRAPH 86)*, 143–150.
- KORN, G. A., AND KORN, T. M. 2000. *Mathematical Handbook for Scientists and Engineers: Definitions, Theorems, and Formulas for Reference and Review*. Courier Dover Publications.
- LI, H., PELLACINI, F., AND TORRANCE, K. 2005. A hybrid monte carlo method for accurate and efficient subsurface scattering. In *Rendering Techniques*, 283–290.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. Graphic.* 22, 3, 759–769.
- NARASIMHAN, S. G., GUPTA, M., DONNER, C., RAMAMOORTHY, R., NAYAR, S., AND JENSEN, H. W. 2006. Acquiring scattering properties of participating media by dilution. *ACM Trans. Graphic.* 25, 1003–1012.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards.
- PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHY, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. Graphic.* 25, 3, 746–753.
- PHARR, M., AND HANRAHAN, P. 2000. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of SIGGRAPH 2000*, 75–84.
- PREMOŽE, S., ASHIKHMIN, M., AND SHIRLEY, P. 2003. Path integration for light transport in volumes. In *Rendering Techniques*, 52–63.
- PREMOZE, S., ASHIKHMIN, M., TESSENDORF, J., RAMAMOORTHY, R., AND NAYAR, S. 2004. Practical rendering of multiple scattering effects in participating media. In *Rendering Techniques*, 363–374.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques*, 41–50.
- TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. Graphic.* 24, 3, 1054–1061.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, 265–272.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, 255–264.
- ZHANG, Z. 1999. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11, 1330–1334.