A Practical Model for Subsurface Light Transport

Henrik Wann Jensen   Stephen R. Marschner   Marc Levoy   Pat Hanrahan

Stanford University

Abstract

This paper introduces a simple model for subsurface light transport in translucent materials. The model enables efficient simulation of effects that BRDF models cannot capture, such as color bleeding within materials and diffusion of light across shadow boundaries. The technique is efficient even for anisotropic, highly scattering media that are expensive to simulate using existing methods. The model combines an exact solution for single scattering with a dipole point source diffusion approximation for multiple scattering. We also have designed a new, rapid image-based measurement technique for determining the optical properties of translucent materials. We validate the model by comparing predicted and measured values and show how the technique can be used to recover the optical properties of a variety of materials, including milk, marble, and skin. Finally, we describe sampling techniques that allow the model to be used within a conventional ray tracer.

Keywords: Subsurface scattering, BSSRDF, reflection models, light transport, diffusion theory, realistic image synthesis

1 Introduction

Accurately modeling the scattering of light by materials is fundamental for realistic image synthesis. Even the most sophisticated light transport algorithms fail to produce convincing results if the local scattering models are too simple. Therefore a great deal of research has gone into describing the scattering of light from materials.

Previous research has focused on developing models for the bidirectional reflectance distribution function (BRDF). The BRDF was introduced by Nicodemus [14] as a simplification of the more general bidirectional surface scattering distribution function (BSSRDF). The BSSRDF can describe light transport between any two rays that hit a surface, whereas the BRDF assumes that light entering a material leaves the material at the same position (Figure 1). This approximation is valid for metals, but it fails for translucent materials, which exhibit significant transport below the surface. Even for many materials that do not seem very translucent, using the BRDF creates a hard, distinctly computer-generated appearance because it does not locally blend surface features such as color and geometry. Only methods that consider subsurface scattering can capture the true appearance of translucent materials, such as marble, cloth, paper, skin, milk, cheese, bread, meat, fruits, plants, fish, ocean water, snow, etc.

1.1 Previous Work

Almost all BRDF models are derived exclusively from surface scattering, with any subsurface scattering approximated by a Lambertian component. An exception is the model by Hanrahan and Krueger [10] which includes an analytic expression for single scattering in a homogeneous, uniformly lit slab. However, all BRDF models ultimately assume that light scatters at one surface point and they do not model subsurface transport from one point to another.

Subsurface transport can be simulated accurately but slowly by solving the full radiative transfer equation [1]. Only a few papers in graphics have taken this approach to subsurface scattering. Dorsey et al. [5] simulated full subsurface scattering using photon mapping to capture the appearance of weathering in stone. Pharr and Hanrahan [15] used scattering functions to simulate subsurface scattering. These approaches, while capable of simulating all of the effects of subsurface scattering, are computationally very expensive compared to the simulation of opaque materials. Techniques based on path sampling are particularly inefficient for highly scattering materials, such as milk and skin, in which light scatters multiple (often several hundred) times before exiting the material. For highly scattering media Stam [17] introduced the use of diffusion theory. He solved a diffusion equation approximation using a multigrid method, and used this method to render clouds with multiple scattering.

Subsurface scattering is also important in medical physics, where models have been developed to describe the scattering of laser light in human tissue [6, 8]. In that context, diffusion theory is often used to predict as well as to measure the optical properties of highly scattering materials. We have extended this theory for use in computer graphics by adding exact single scattering, support for arbitrary geometry, and a practical sampling technique for rendering.

In measurements of appearance for computer graphics, subsurface scattering has rarely been considered. Debevec et al. [3] measured light reflection from human faces, which included contributions from subsurface scattering, but they did not relate the data to the physical properties of the material. Again building on medical physics research [8, 9], we have extended a methodology developed for measuring biological tissues into a rapid image-based appearance measurement technique for translucent materials. This method examines the radial reflectance profile resulting from a beam illuminating the sample material. By fitting an expression derived from diffusion theory it is possible to estimate the absorption and scattering properties of the material.

2 Theory

The BSSRDF, S, relates the outgoing radiance, $L_o(x_o, \vec{\omega}_o)$ at the point x_o in direction $\vec{\omega}_o$, to the incident flux, $\Phi_i(x_i, \vec{\omega}_i)$ at the point x_i from direction $\vec{\omega}_i$ [14]:

$$dL_o(x_o, \vec{\omega}_o) = S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o)\, d\Phi_i(x_i, \vec{\omega}_i).$$

The BRDF is an approximation of the BSSRDF for which it is assumed that light enters and leaves at the same point (i.e., $x_o = x_i$). Given a BSSRDF, the outgoing radiance is computed
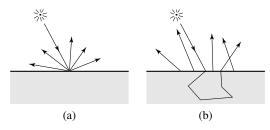
Figure 1: Scattering of light in (a) a BRDF, and (b) a BSSRDF.

by integrating the incident radiance over incoming directions *and area*, A:

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

Light propagation in a participating medium is described by the radiative transport equation, often referred to in computer graphics as the volume rendering equation:

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

In this equation, the properties of the medium are described by the absorption coefficient σ_a, the scattering coefficient σ_s, and the phase function $p(\vec{\omega}, \vec{\omega}')$. The extinction coefficient σ_t is defined as, $\sigma_t = \sigma_a + \sigma_s$. We assume the phase function is normalized, $\int_{4\pi} p(\vec{\omega}, \vec{\omega}') \, d\omega' = 1$ and is a function only of the phase angle, $p(\vec{\omega}, \vec{\omega}') = p(\vec{\omega} \cdot \vec{\omega}')$. The mean cosine, g, of the scattering angle is

$$g = \int_{4\pi} (\vec{\omega} \cdot \vec{\omega}') p(\vec{\omega} \cdot \vec{\omega}') \, d\omega'.$$

If g is positive, the phase function is predominantly forward scattering; if g is negative, backward scattering dominates. A constant phase function results in isotropic scattering ($g = 0$).

For an infinitesimal beam entering a homogeneous medium, the incoming radiance will decrease exponentially with distance s. This is referred to as the *reduced intensity*:

$$L_{ri}(x_i + s\vec{\omega}_i, \vec{\omega}_i) = e^{-\sigma_t s} L_i(x_i, \vec{\omega}_i).$$

The first-order scattering of the reduced intensity, L_{ri}, may be treated as a volumetric source:

$$Q(x, \vec{\omega}) = \sigma_s \int_{4\pi} p(\vec{\omega}', \vec{\omega}) L_{ri}(x, \vec{\omega}') \, d\omega'.$$

To gain insight into the volumetric behavior of light propagation, it is useful to integrate the radiative transport equation over all directions $\vec{\omega}$ at a point x which yields

$$\vec{\nabla} \cdot \vec{E}(x) = -\sigma_a \phi(x) + Q_0(x). \tag{1}$$

This equation relates the scalar irradiance, or fluence, $\phi(x) = \int_{4\pi} L(x, \vec{\omega}) \, d\omega$, and the vector irradiance, $\vec{E}(x) = \int_{4\pi} L(x, \vec{\omega}) \vec{\omega} \, d\omega$. In the absence of loss due to absorption or gain due to a volumetric light source ($Q_0 = 0$), the divergence of the vector irradiance equals zero. In this equation, we introduce a 0th-order source term, Q_0, and later we will need the 1st-order source term, $\vec{Q}_1$, where

$$Q_0(x) = \int_{4\pi} Q(x, \vec{\omega}) \, d\omega, \quad \vec{Q}_1(x) = \int_{4\pi} Q(x, \vec{\omega}) \vec{\omega} \, d\omega.$$

S	BSSRDF
R_d	Diffuse BSSRDF
F_r	Fresnel reflectance
F_t	Fresnel transmittance
F_{dr}	Diffuse Fresnel reflectance
$\vec{E}$	Vector irradiance
ϕ	Radiant fluence
σ_a	Absorption coefficient
σ_s	Scattering coefficient
σ_t	Extinction coefficient
σ_t'	Reduced extinction coefficient
σ_{tr}	Effective extinction coefficient
D	Diffusion constant
α	Albedo
p	Phase function
η	Relative index of refraction
g	Mean cosine of the scattering angle
Q	Volume source distribution
Q_0	0th-order source distribution
$\vec{Q}_1$	1st-order source distribution

Figure 2: Selected symbols.

2.1 The Diffusion Approximation

The diffusion approximation is based on the observation that the light distribution in highly scattering media tends to become isotropic. This is true even if the initial light source distribution and the phase function are highly anisotropic. Each scattering event blurs the light distribution, and as a result the light distribution tends toward uniformity as the number of scattering events increases.

In this situation, the radiance may be approximated by a two-term expansion involving the radiant fluence and the vector irradiance:

$$L(x, \vec{\omega}) = \frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \vec{\omega} \cdot \vec{E}(x).$$

The constants are determined by the definitions of fluence and vector irradiance.

The diffusion equation follows from this approximation. Specifically, we substitute this two-term expansion of the radiance into the radiative transport equation and then integrate over $\vec{\omega}$; for the algebraic details consult Ishimaru [12]. The result is

$$\vec{\nabla} \phi(x) = -3\sigma_t' \vec{E}(x) + \vec{Q}_1(x). \tag{2}$$

Here we have used the reduced extinction coefficient, σ_t', which is given by

$$\sigma_t' = \sigma_s' + \sigma_a \quad \text{where} \quad \sigma_s' = \sigma_s(1 - g).$$

The reduced scattering coefficient σ_s' scales the original scattering coefficient by a factor of $(1 - g)$. Intuitively, once light becomes isotropic, only backward scattering terms change the net flux; forward scattering is indistinguishable from no scattering.

In the case where there are no sources, or where the sources are isotropic, $\vec{Q}_1$ vanishes from Equation 2. Then the vector irradiance is the gradient of the scalar fluence,

$$\vec{E}(x) = -D\vec{\nabla}\phi(x).$$

Here $D = \frac{1}{3\sigma_t'}$ is the diffusion constant. This equation makes precise the intuitive notion that there is net energy flow (i.e., non-zero vector irradiance) from regions of high energy density (high fluence) to regions of low energy density.

Finally, substituting Equation 2 into Equation 1, we arrive at the classic diffusion equation

$$D\nabla^2 \phi(x) = \sigma_a \phi(x) - Q_0(x) + 3D\vec{\nabla} \cdot \vec{Q}_1(x).$$

The diffusion equation has a simple solution in the case of a single isotropic point light source in an infinite medium.

$$\phi(x) = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr} r(x)}}{r(x)},$$

where Φ is the power of the point light source, r is the distance to the location of the point source, and $\sigma_{tr} = \sqrt{3\sigma_a \sigma'_t}$ is the effective transport coefficient. The point source results in an energy density in the volume with an exponential falloff.

In the case of a scattering medium in a finite region of space, the diffusion equation must be solved subject to the appropriate boundary conditions. The boundary condition is that the net inward diffuse flux is zero at each point, x_s, on the surface

$$\int_{2\pi_-} L(x_s, \vec{\omega})(\vec{\omega} \cdot \vec{n}(x_s)) \, d\omega = 0.$$

Here, $2\pi_-$ denotes integration over the hemisphere of inward directions. Using the two-term expansion, the boundary condition is

$$\phi(x_s) - 2D(\vec{n} \cdot \vec{\nabla})\phi(x_s) = 0. \tag{3}$$

The minus sign in the second term results from the convention that the surface normal points outward, whereas the integral is over inward directions.

Equation 3 covers the case where the two layers have matching indices of refraction, but another important case is where these indices differ. When an interface exists between media with different refractive indices, there is a reflection at the interface. Assuming F_r is the Fresnel formula for the reflectance at a dielectric interface, the average diffuse Fresnel reflectance is

$$F_{dr} = \int_{2\pi} F_r(\eta, \vec{n} \cdot \vec{\omega}')(\vec{n} \cdot \vec{\omega}') \, d\omega',$$

where η is the relative index of refraction of the medium with the reflected ray to the other medium. F_{dr} may be computed analytically from the Fresnel formula [13]. However, we will use a rational approximation of the measured diffuse reflectance [7]:

$$F_{dr} = -\frac{1.440}{\eta^2} + \frac{0.710}{\eta} + 0.668 + 0.0636\eta.$$

The resulting boundary condition between two media with different indices of refraction is

$$\int_{2\pi_-} L(x, \vec{\omega})(\vec{\omega} \cdot \vec{n}_-) d\omega = F_{dr} \int_{2\pi_+} L(x, \vec{\omega})(\vec{\omega} \cdot \vec{n}_+) \, d\omega.$$

Here the $+$ and $-$ subscript means outward and inward directions respectively. This yields

$$\phi(x_s) - 2D(\vec{n} \cdot \vec{\nabla})\phi(x_s) = F_{dr} \left[\phi(x_s) + 2D(\vec{n} \cdot \vec{\nabla})\phi(x_s) \right].$$

Note that the difference in signs between the two sides of this equation occurs because one integral is over outward directions and the other is over inward directions. Rearranging terms,

$$\phi(x_s) - 2AD(\vec{n} \cdot \vec{\nabla})\phi(x_s) = 0.$$

This boundary condition is the same as when the indices of refraction match (Equation 3); the only difference is that $2D$ is replaced by $2AD$, where

$$A = \frac{1 + F_{dr}}{1 - F_{dr}}.$$

Finally, the boundary condition allows us to compute the diffuse BSSRDF, R_d. R_d is equal to the radiant exitance divided by the incident flux. The radiant exitance leaving the surface $(\vec{n} \cdot \vec{E}(x_s))$ is equal to the gradient of the fluence at the surface

$$R_d(r) = -D \frac{(\vec{n} \cdot \vec{\nabla}\phi)(x_s)}{d\Phi_i(x_i)},$$

where $r = \|x_s - x_i\|$.

In the case of finite media, the diffusion equation does not in general have an analytical solution. In this paper we are interested in subsurface reflection, which is often modeled as a semi-infinite plane-parallel medium. Several authors have analyzed the plane-parallel problem for simple source geometries, in particular, approximations of a cylindrical beam entering the media. Exact formulas exist, but they involve an infinite sum of Bessel functions [9, 16]. We seek a simple formula suitable for modeling subsurface reflection that does not involve infinite sums or numerical solution of a partial differential equation.

Eason [6] and Farrell et al. [8] have developed a method for approximating the volumetric source distribution using two point sources; that is, a dipole. Eason introduced this idea and derived explicit formulae for the dipoles for various source geometries, such as a cylindrical beam, by expanding the source distributions in terms of their moments. Farrell et al. proposed using a single dipole to represent the incident source distribution. They found a single dipole to be as accurate as, or, in some cases, more accurate than using the diffusion approximation with the true source distribution.

The dipole method consists of positioning two point sources near the surface in such a way as to satisfy the required boundary condition [6] (see Figure 3). One point source, the positive real light source, is located at the distance z_r beneath the surface, and the other, the negative virtual light source, is located above the surface at a distance $z_v = z_r + 4AD$. The resulting fluence is

$$\phi(x) = \frac{\Phi}{4\pi D} \left(\frac{e^{-\sigma_{tr} d_r}}{d_r} - \frac{e^{-\sigma_{tr} d_v}}{d_v} \right),$$

where $d_r = \|x - x_r\|$ is the distance from x to the real source, and $d_v = \|x - x_v\|$ is the distance from x to the virtual source. Farrell et al. [8] proposed positioning the real light source at distance $z_r = 1/\sigma'_t$, or one mean free path, below the surface. They only considered light parallel to the normal. For other light directions reciprocity can be enforced by still placing the light source $1/\sigma'_t$ straight below x_i.

The diffuse reflectance due to the dipole source can now be computed.

$$R_d(r) = -D \frac{(\vec{n} \cdot \vec{\nabla}\phi(x_s))}{d\Phi_i}$$
$$= \frac{\alpha'}{4\pi} \left[(\sigma_{tr} d_r + 1) \frac{e^{-\sigma_{tr} d_r}}{\sigma'_t d_r^3} + z_v (\sigma_{tr} d_v + 1) \frac{e^{-\sigma_{tr} d_v}}{\sigma'_t d_v^3} \right]. \tag{4}$$

Lastly, we need to take into account the Fresnel reflection at the boundary for both the incoming light and the outgoing radiance.

$$S_d(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) = \frac{1}{\pi} F_t(\eta, \vec{\omega}_i) R_d(\|x_i - x_o\|) F_t(\eta, \vec{\omega}_o) \tag{5}$$

where S_d is the diffusion term of the BSSRDF. This term represents multiple scattering (one scattering event is already included in the conversion to a point source). The next section explains how to compute the contribution due to single scattering.
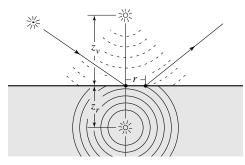
Figure 3: An incoming ray is transformed into a dipole source for the diffusion approximation.

2.2 Single Scattering Term

Hanrahan and Krueger [10] have derived a BRDF model for subsurface reflection that analytically computes the total first-order scattering from a flat, uniformly lit, homogeneous slab. In this section, we show how their BRDF can be extended to a BSSRDF in order to account for local variations in lighting over the surface.

The total outgoing radiance, $L_o^{(1)}$, due to single scattering is computed by integrating the incident radiance along the refracted outgoing ray (see Figure 4):

$$L_o^{(1)}(x_o, \vec{\omega}_o) = \sigma_s(x_o) \int_{2\pi} F \, p(\vec{\omega}_i' \cdot \vec{\omega}_o') \int_0^\infty e^{-\sigma_{tc} s} L_i(x_i, \vec{\omega}_i) \, ds \, d\vec{\omega}_i \quad (6)$$

$$= \int_A \int_{2\pi} S^{(1)}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) \, L_i(x_i, \vec{\omega}_i) \, (\vec{n} \cdot \vec{\omega}_i) \, d\omega_i dA(x_i).$$

Here $F = F_t(\eta, \vec{\omega}_o) F_t(\eta, \vec{\omega}_i)$ is the product of the two Fresnel transmission terms, and $\vec{\omega}_i'$ and $\vec{\omega}_o'$ are the refracted incoming and outgoing directions. The combined extinction coefficient σ_{tc} is given by $\sigma_{tc} = \sigma_t(x_o) + G\sigma_t(x_i)$, where G is a geometry factor; for a flat surface $G = \frac{|\vec{n}_i \cdot \vec{\omega}_o'|}{|\vec{n}_i \cdot \vec{\omega}_i'|}$. The single scattering BSSRDF, $S^{(1)}$, is defined implicitly by the second line of this equation. Note that there is a change of variables between the first line, which integrates only over the configurations where the two refracted rays intersect, and the second line, which integrates over all incoming and outgoing rays. This implies that the distribution $S^{(1)}$ contains a delta function.

2.3 The BSSRDF Model

The complete BSSRDF model is a sum of the diffusion approximation and the single scattering term:

$$S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) = S_d(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) + S^{(1)}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o)$$

Here S_d is evaluated using Equation 5 and $S^{(1)}$ is evaluated using Equation 6. The parameters for the BSSRDF are: σ_a, σ_s', η, and possibly a phase function (without a phase function the scattering can be modeled as isotropic). This model accounts for light transport between different locations on the surface, and it simulates both the directional component (due to single scattering) as well as the diffuse component (due to multiple scattering).

Finally, note the distances involved in both the single scattering term and the diffusion approximations. The average exit point is approximately one mean free path from the entry point. However, these two mean free paths have quite different length scales. In the single scattering case, the mean free path equals $1/\sigma_t$; in the diffusion case, the mean free path equals $1/\sigma_{tr}$. For translucent materials where $\sigma_a \ll \sigma_s'$ and consequently $\sigma_{tr} \ll \sigma_t$, the single scattering term decreases much faster than the diffusion term as the distance to x_o increases.
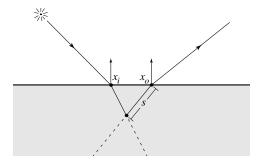

Figure 4: Single scattering occurs only when the refracted incoming and outgoing rays intersect, and is computed as an integral over path length s along the refracted outgoing ray.

2.4 BRDF Approximation

We can approximate the BSSRDF with a BRDF by assuming that the incident illumination is uniform. This assumption makes it possible to integrate the BSSRDF over the surface. By integrating the diffusion term we find the total diffuse reflectance R_d of the material as:

$$R_d = 2\pi \int_0^\infty R_d(r) \, r \, dr = \frac{\alpha'}{2} \left(1 + e^{-\frac{4}{3} A \sqrt{3(1-\alpha')}} \right) e^{-\sqrt{3(1-\alpha')}} \, .$$

Notice how the diffuse reflectance only depends on the reduced albedo and the internal reflection parameter A.

The integration of the single scattering term results in the model presented in [10]. For a semi-infinite medium this gives:

$$f_r^{(1)}(x, \vec{\omega}_i, \vec{\omega}_o) = \alpha F \frac{p(\vec{\omega}_i' \cdot \vec{\omega}_o')}{|\vec{n} \cdot \vec{\omega}_i'| + |\vec{n} \cdot \vec{\omega}_o'|} \, .$$

The complete BRDF model is the sum of the diffuse reflectance scaled by the Fresnel term and the single scattering approximation:
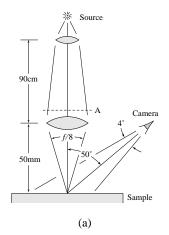
$$f_r(x, \vec{\omega}_i, \vec{\omega}_o) = f_r^{(1)}(x, \vec{\omega}_i, \vec{\omega}_o) + F \frac{R_d}{\pi} \, .$$

This model has the same parameters as the BSSRDF. It is similar to the BRDF model presented in [10], but with the important difference that the amount of diffusely reflected light is computed from the intrinsic material parameters. The BRDF approximation is useful for opaque materials, which have a very short mean free path.

3 Measuring the BSSRDF

To verify our BSSRDF model, and to determine appropriate parameters for rendering different kinds of materials, we used the diffusion theory of Section 2 to make measurements of subsurface scattering in several media. Our measurement approach applies to translucent materials for which $\sigma_a \ll \sigma_s$, implying that far enough away from the point of illumination, we may neglect single scattering and use the diffusion term to relate measurements to material parameters.

When multiple scattering dominates, Equation 4 predicts the radiant exitance per unit incident flux that will be observed due to a narrow incident beam, as a function of distance from the point of illumination. To make the corresponding measurement, we illuminate the surface of a sample with a tightly focused beam of white light and take a photograph using a 3-CCD video camera to observe the radiant exitance across the entire surface. We keep our observations at constant angles so that the Fresnel term remains constant for all the measurements. Figure 5(a) illustrates our measurement setup.

Material	σ'_s [mm^{-1}]			σ_a [mm^{-1}]			Diffuse Reflectance			η
	R	G	B	R	G	B	R	G	B	
Apple	2.29	2.39	1.97	0.0030	0.0034	0.046	0.85	0.84	0.53	1.3
Chicken1	0.15	0.21	0.38	0.015	0.077	0.19	0.31	0.15	0.10	1.3
Chicken2	0.19	0.25	0.32	0.018	0.088	0.20	0.32	0.16	0.10	1.3
Cream	7.38	5.47	3.15	0.0002	0.0028	0.0163	0.98	0.90	0.73	1.3
Ketchup	0.18	0.07	0.03	0.061	0.97	1.45	0.16	0.01	0.00	1.3
Marble	2.19	2.62	3.00	0.0021	0.0041	0.0071	0.83	0.79	0.75	1.5
Potato	0.68	0.70	0.55	0.0024	0.0090	0.12	0.77	0.62	0.21	1.3
Skimmilk	0.70	1.22	1.90	0.0014	0.0025	0.0142	0.81	0.81	0.69	1.3
Skin1	0.74	0.88	1.01	0.032	0.17	0.48	0.44	0.22	0.13	1.3
Skin2	1.09	1.59	1.79	0.013	0.070	0.145	0.63	0.44	0.34	1.3
Spectralon	11.6	20.4	14.9	0.00	0.00	0.00	1.00	1.00	1.00	1.3
Wholemilk	2.55	3.21	3.77	0.0011	0.0024	0.014	0.91	0.88	0.76	1.3

(a)   (b)

Figure 5: (a) Measurement apparatus, (b) measured parameters for several materials.

Because the signal falls off exponentially away from the point of illumination, the measurement must span a wide dynamic range. To this end we used a series of different exposure times, ranging from 1 millisecond to 4 seconds, and assembled a high-dynamic-range image using a modified version of Debevec and Malik's technique [4]. To reduce the effects of stray light and fixed-pattern CCD noise, we subtracted a dark image, taken with the illumination beam blocked just before the focusing lens (point A in Figure 5(a)), from each measurement and reference image. The resulting images had a dynamic range of around 10^5 (the small amount of total energy in the image reduces the effects of lens and camera flare, allowing higher dynamic range than might otherwise be possible).

To interpret the measurements, we examined only a 1D slice of each measurement image, corresponding to a line on the surface through the illumination point and perpendicular to the camera's view direction. Under the assumption that light exits diffusely[1], the pixel values p_i in this slice (see Figure 6 for an example) are measurements of radiant exitance as a function of distance on the surface. Since R_d gives the ratio of this quantity to Φ, $p_i = K\Phi R_d(r_i)$, where K is an unknown constant. To eliminate the scale factor, we also took a reference image with the sample replaced by a white ideal diffuse reflector (Labsphere Spectralon, reflectance > 0.99). By summing all the pixels in this image, we can integrate the radiant exitance to get the total flux exiting the surface, which for this special material is equal to the incident flux Φ. With the same constant K as above, this sum is $K\Phi/A$, where A is the (known) area on the sample's surface subtended by one pixel. The measured value for $R_d(r_i)$ can then be computed as $p_i/(K\Phi)$.

In principle, σ_a and σ'_s can be determined by fitting the relative reflectance curve with Equation 4 over a range of distances far enough from the illumination point to allow the use of diffusion theory [8]. However, we found this fitting problem to be ill-conditioned enough that the uncertainty in the resulting parameters led to too much uncertainty in the appearance of the material, especially the total diffuse reflectance.

We remove this ill-conditioning by measuring the total diffuse reflectance R (which is the sum of the measurement image divided by the sum of the reference image) and computing the least-squares fit subject to the constraint $\int R_d \, dA = R$.

Figure 6 shows how these measurements confirm the diffusion theory for a sample of white marble (only the camera's green chan-

---

[1] We verified this assumption for marble by examining the reflectance for different outgoing angles, and it closely resembled a Lambertian material scaled by a Fresnel transmission term.
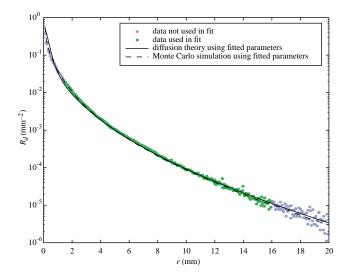

Figure 6: Measurements for marble (green wavelength band) plotted with fit to diffusion theory and confirming Monte Carlo simulation.

nel is shown). Fitting the theory (solid line) to the data (points) led to the parameters $\sigma_a = 0.0041$/mm, $\sigma'_s = 2.6$/mm. The reflectance computed by a Monte Carlo simulation using these values (dashed line) confirms the correctness of the computed parameters. Fitted values for several other materials appear in the table in Figure 5(b). Note, that we used empirical values for the index of refraction for most of the materials. Also note that the diffusion theory is assuming that $\sigma_s \gg \sigma_a$, and as such the parameters for the relatively opaque materials (such as the blue wavelength in ketchup) may be less accurate.

4 Rendering Using the BSSRDF

The BSSRDF model derived in the theory section only applies to semi-infinite homogeneous media. A similar derivation is not possible in the presence of arbitrary geometry and texture variation. However, we can use some of the intuition behind the theory to extend it to a practical model for computer graphics. Specifically, we need to consider:
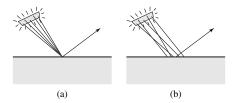
(a)   (b)

Figure 7: (a) Sampling a BRDF (traditional sampling), (b) sampling a BSSRDF (the sample points are distributed both over the surface as well as the light).

- Efficient integration of the BSSRDF including importance sampling

- Single scattering evaluation for arbitrary geometry

- Diffusion approximation for arbitrary geometry

- Texture (spatial variation on the object surface).

In this section we explain how to do this in a ray-tracing context.

Integrating the BSSRDF: At each ray-object intersection traditional lighting models (based on BRDFs) need just a point and a normal to compute the outgoing radiance (Figure 7(a)). For the BSSRDF it is necessary to integrate the incoming lighting over an area of the surface (Figure 7(b)). We do this by stochastically sampling the location of *both* endpoints of the shadow ray — this can be seen as an extension of the classical distribution ray tracing technique for sampling area light sources [2]. To efficiently sample locations on the surface we exploit the exponential falloff in the diffusion term and the single scattering term. We sample the two terms of the BSSRDF separately, since the single scattering sample locations must be along the refracted outgoing ray whereas the diffusion samples should be distributed around x_o.

More specifically, for the diffusion term, we use standard Monte Carlo techniques to randomly sample the surface with density $(\sigma_{tr} e^{-\sigma_{tr} d})$ at some distance d from x_o.

Single scattering is reparameterized since the incoming ray and the outgoing ray must intersect. Our technique is explained in the following section.

Single scattering evaluation for arbitrary geometry: Single scattering is evaluated using Monte Carlo integration along the refracted outgoing ray. We pick a random distance, $s'_o = \log(\xi)/\sigma_t(x_o)$, along the refracted outgoing ray. Here $\xi \in]0,1]$ is a uniformly distributed random number. For this sample location we compute the outscattered radiance as:

$$L_o^{(1)}(x_o, \vec{\omega}_o) = \frac{\sigma_s(x_o) F p(\vec{\omega}_i \cdot \vec{\omega}_o)}{\sigma_{tc}} e^{-s'_i \sigma_t(x_i)} e^{-s'_o \sigma_t(x_o)} L_i(x_i, \vec{\omega}_i).$$

Here s'_i is the distance that the sample ray moves through the material. Optimizing this equation to sample direct illumination (with shadow rays) is difficult for arbitrary geometry since it requires finding the point at the surface where the shadow ray is refracted. However, in practice a good approximation can be found by using a shadow ray that does not refract at the surface — this assumes that the light source is far away compared to the mean free path of the medium. We can use Snell's law to estimate the true refracted distance through the medium of the incoming ray:

$$s'_i = s_i \frac{|\vec{\omega}_i \cdot \vec{n}_i|}{\sqrt{1 - \left(\frac{1}{\eta}\right)^2 (1 - |\vec{\omega}_i \cdot \vec{n}(x_i)|^2)}} .$$

Here s_i is the observed distance and s'_i is the refracted distance.
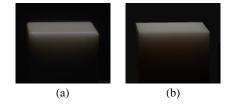

(a)   (b)

Figure 8: Scattering of laser light in a marble block. The marble block is 40mm. wide and has a significant amount of subsurface scattering. The picture on the left is a photograph of the marble block, and the picture on the right is a synthetic rendering of a similarly sized cube using the BSSRDF model and the measured scattering properties of the marble. Note how the appearance of the two images is very similar.

Diffusion approximation for arbitrary geometry: An important component of the diffusion approximation is the use of the dipole source. If the geometry is locally flat we can get a very good approximation by using a similar dipole source configuration as that for flat materials (i.e., we always place the light source $1/\sigma'_t$ straight below x_i). Special care must be taken in the presence of highly curved surfaces; we handle this case by always evaluating the diffusion term with a minimum distance of $1/\sigma'_t$. In this way we eliminate singularities at sharp edges where the source can be placed arbitrarily close to x_o. We found this approach to work very well in our experiments.

Texture: We approximate textured materials by making a few small changes to the usage of the BSSRDF. We only consider texture variation at the surface — effects due to volumetric texture variation would require a full participating media simulation. For the diffusion approximation we always use the material parameters at x_i, which ensures a natural local blending of the texture properties. For the single scattering term we use $\sigma_s(x_o)$ and $\sigma_t(x_o)$ along the refracted outgoing ray, and $\sigma_t(x_i)$ along the refracted incident ray. This variation is included in Equation 6.

5   Results

We have implemented the BSSRDF model in a Monte Carlo ray tracer, and in this section we will present a number of experimental results obtained with this implementation. All simulations have been done on a dual 800MHz Pentium III PC running Linux and the images have been rendered with 4 samples per pixel and a width of 1024 pixels.

Our first simulation is shown in Figure 8, which compares a side photograph of a marble cube illuminated from above with a synthetic rendering. The synthetic image is rendered using the BSSRDF model and the measured parameters for marble (from the table in Figure 5). We only used a simple cube to approximate the rounded marble block, so there are natural visible differences along the edges. Nonetheless, the BSSRDF model faithfully renders the appearance including the scattered light exiting from the side of the marble cube.

Figure 9 shows several different simulations of subsurface scattering in a marble bust (1.3 million triangles) illuminated from behind. The BSSRDF simulation mostly matches the appearance of the full Monte Carlo simulation, yet is significantly faster (5 minutes vs. 1250 minutes). The hair at the back of the head is slightly darker in the BSSRDF simulation; we believe this is due to the forced $1/\sigma'_t$ distance in the diffusion approximation. A similar rendering was done using photon mapping in [5] in roughly 12 minutes (scaled to the speed of our computer). However, the photon mapping method requires a full 3D-description of the material, it requires memory to store the photons, and it becomes costly for
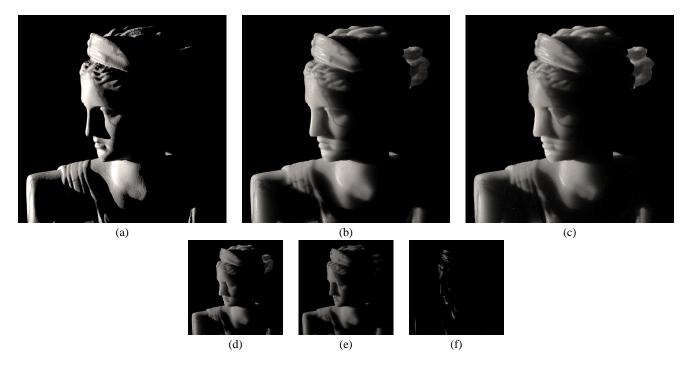
Figure 9: A simulation of subsurface scattering in a marble bust. The marble bust is illuminated from behind and rendered using: (a) the BRDF approximation (in 2 minutes), (b) the BSSRDF approximation (in 5 minutes), and (c) a full Monte Carlo simulation (in 1250 minutes). Notice how the BSSRDF model matches the appearance of the Monte Carlo simulation, yet is significantly faster. The images in (d–f) show the different components of the BSSRDF: (d) single scattering term, (e) diffusion term, and (f) Fresnel term.

highly scattering materials (such as milk and skin).

A particularly interesting aspect of the BSSRDF simulation is that it is able to capture the smooth appearance of the marble surface. In comparison the BRDF simulation gives a very hard appearance where even tiny bumps on the surface are visible (this is a classic problem in realistic image synthesis where objects often look hard and unreal).

For the marble we used synthetic scattering and absorption coefficients, since we wanted to test the difficult case when the average scattering albedo is 0.5 (here the contribution from diffusion and single scattering is approximately the same). Figure 9 demonstrates how the sum of both single scattering and the diffusion term is necessary to match the Monte Carlo simulation.

Figure 10 contains three renderings of milk. The first rendering uses a diffuse reflection model; the others use the BSSRDF model and our measurements for skim milk and whole milk. Notice how the diffuse milk looks unreal and too opaque compared to the BSSRDF images, even though multiple scattering dominates and the radiant exitance due to subsurface scattering is very diffuse. It is interesting that the BSSRDF simulations are capable of capturing the subtle details in the appearance of milk, making the milk look more bluish at the front and more reddish at the back. This is due to Rayleigh scattering that causes shorter wavelengths of light to be scattered more than longer wavelengths.
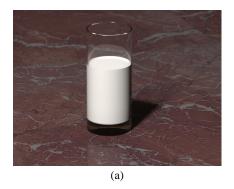
Skin is a material that is particularly difficult to render using methods that simulate subsurface scattering by sampling ray paths through the material. This is due to the fact that skin is highly scattering (typical albedo is 0.95) and also very anisotropic (typical average cosine of the scattering angle is 0.85). Both of these properties mean that the average number of scattering events of a photon is very high (often more than 100). In addition skin is very translucent, and it cannot be rendered correctly using a BRDF (see Figure 11). A complete skin model requires multiple layers, but a
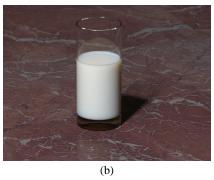
reasonable approximation can be obtained using just one layer. In Figure 11 we have rendered a simple face model using the BSSRDF and our measured values for skin (skin1). Here we also used the Henyey-Greenstein phase function [11] with $g = 0.85$ as the estimated mean cosine of the scattering angle. The skin measurements are from an arm (which is likely more translucent than skin on the face), but the overall appearance is still realistic considering the lack of spatial variation (texture). The BSSRDF gives the skin a soft appearance, and it renders the color bleeding in the shadow region below the nose. Here, the absorption by blood is particularly noticeable as the light that scatters deep in the skin is redder. For this simulation the diffusion term is much larger than the single scattering term. This means that skin reflects light fairly diffusely, but also that internal color bleeding is an important factor. The BRDF image was rendered in 7 minutes, the BSSRDF image was rendered in 17 minutes.

6 Conclusion and Future Work

In this paper we have presented a new practical BSSRDF model for computer graphics. The model combines a dipole diffusion approximation with an accurate single scattering computation. We have shown how the model can be used to measure the scattering properties of translucent materials, and how the measured values can be used to reproduce the results of the measurements as well as synthetic renderings. We evaluate the BSSRDF by sampling the incoming light over the surface, and we demonstrate how this technique is capable of capturing the soft and smooth appearance of translucent materials.

In the future we plan to extend the model to multiple layers as well as include support for efficient global illumination.

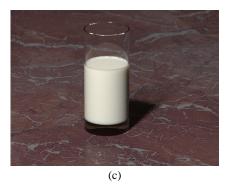(a)   (b)   (c)

Figure 10: A glass of milk: (a) diffuse (BRDF), (b) skim (BSSRDF) and (c) whole (BSSRDF). (b) and (c) are using our measured values. The rendering times are 2 minutes for (a), and 4 minutes for (b) and (c); this includes caustics and global illumination on the marble table and a depth-of-field simulation.

BRDF

BSSRDF

Figure 11: A face rendered using the BRDF model (top) and the BSSRDF model (bottom). We used our measured values for skin (skin1) and the same lighting conditions in both images (the BRDF image also includes global illumination). The face geometry has been modeled by hand; the lip-bumpmap is handpainted, and the bumpmap on the skin is based on a gray-scale macro photograph of a piece of skin. Even with global illumination the BRDF gives a hard appearance. Compare this to the faithful soft appearance of the skin in the BSSRDF simulation. In addition the BSSRDF captures the internal color bleeding in the shadow region under the nose.

7 Acknowledgements

References

[1] S. Chandrasekhar. *Radiative Transfer*. Oxford Univ. Press, 1960.

[2] R. L. Cook, T. Porter, and L. Carpenter. Distributed ray tracing. In *ACM Computer Graphics (SIGGRAPH'84)*, volume 18, pages 137–145, July 1984.

[3] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Computer Graphics Proceedings, Annual Conference Series, 2000*, pages 145–156, July 2000.

[4] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics Proceedings, Annual Conference Series, 1997*, pages 369–378, August 1997.

[5] J. Dorsey, A. Edelman, H. W. Jensen, J. Legakis, and H. K. Pedersen. Modeling and rendering of weathered stone. In *Computer Graphics Proceedings, Annual Conference Series, 1999*, pages 225–234, August 1999.

[6] G. Eason, A. Veitch, R. Nisbet, and F. Turnbull. The theory of the backscattering of light by blood. *J. Physics*, 11:1463–1479, 1978.

[7] W. G. Egan and T. W. Hilgeman. *Optical Properties of Inhomogeneous Materials*. Academic Press, New York, 1979.

[8] T. J. Farell, M. S. Patterson, and B. Wilson. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Med. Phys.*, 19:879–888, 1992.

[9] R. A. Groenhuis, H. A. Ferwerda, and J. J. Ten Bosch. Scattering and absorption of turbid materials determined from reflection measurements. 1: Theory. *Applied Optics*, 22:2456–2462, 1983.

[10] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. In *ACM Computer Graphics (SIGGRAPH'93)*, pages 165–174, August 1993.

[11] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the galaxy. *Astrophysics Journal*, 93:70–83, 1941.

[12] A. Ishimaru. *Wave Propagation and Scattering in Random Media*, volume 1. Academic Press, New York, 1978.

[13] G. Kortum. *Reflectance Spectroscopy*. Springer-Verlag, 1969.

[14] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. Monograph 161, National Bureau of Standards (US), October 1977.

[15] M. Pharr and P. Hanrahan. Monte Carlo evaluation of non-linear scattering equations for subsurface reflection. In *Computer Graphics Proceedings, Annual Conference Series, 2000*, pages 75–84, July 2000.

[16] L. Reynolds, C. Johnson, and A. Ishimaru. *Applied Optics*, 15:2059, 1976.

[17] J. Stam. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 1995*. Eurographics, June 1995.