

Directional Dipole for Subsurface Scattering in Translucent Materials

Jeppe Revall Frisvad¹

Toshiya Hachisuka²

Thomas Kim Kjeldsen³

¹Technical University of Denmark

²Aarhus University

³The Alexandra Institute

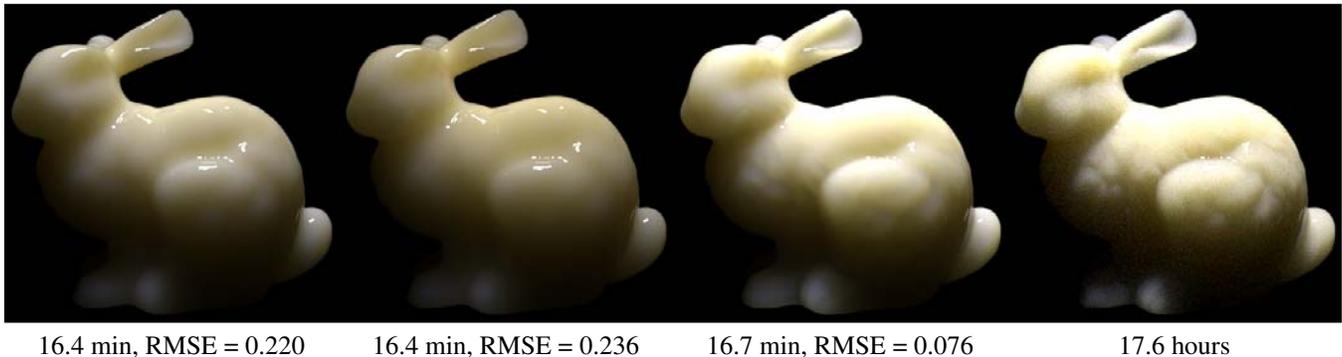


Figure 1: Stanford bunny made of white grapefruit juice [Narasimhan et al. 2006] and rendered using, from left to right, the standard dipole [Jensen et al. 2001], the better dipole [d’Eon 2012; d’Eon and Irving 2011], our directional dipole, and unbiased path tracing [Rushmeier 1988]. Our model captures translucency effects that are present in path tracing but not in other analytical BSSRDF models.

Abstract

Rendering translucent materials using Monte Carlo ray tracing is a challenge. For such materials, rendering time often becomes prohibitively long due to the large number of subsurface scattering events. A faster approach is to use an analytical model derived from diffusion theory. However, the efficiency of the analytical models comes at the cost of losing translucency effects in the rendered result. We present a new analytical model for subsurface scattering. Our model captures translucency effects that are present in the reference images but remain absent with existing analytical models. The key difference is that our model is based on ray source diffusion, rather than point source diffusion. A ray source corresponds better to the light that refracts through the surface of a translucent material. Using this ray source, we are able to take the direction of the incident light ray and the direction toward the point of emergence into account. We use a dipole construction similar to that of the standard dipole model, but we now have positive and negative ray sources instead of point sources. Our model is as efficient as previous dipole models while the rendered images are significantly closer to the references. Unlike some previous work, our model is fully analytic and requires no precomputation.

1 Introduction

The rendering of translucent materials, such as skin, foods, stones, and many other natural materials, has many important use cases in computer graphics. For more than a decade, the dipole approximation for subsurface scattering [Jensen et al. 2001] has proven to be a fast practical way of rendering such materials. This standard dipole model is however built upon a number of assumptions which are often violated. One significant assumption is that incident light is directionally uniform. This assumption is evident from the fact that the standard dipole model is a function only of the distance between the point of incidence and the point of emergence. The same assumption has been used in improved analytical models which were introduced recently [d’Eon and Irving 2011; d’Eon 2012]. However, the lighting distribution in realistic scene configurations is rarely directionally uniform.

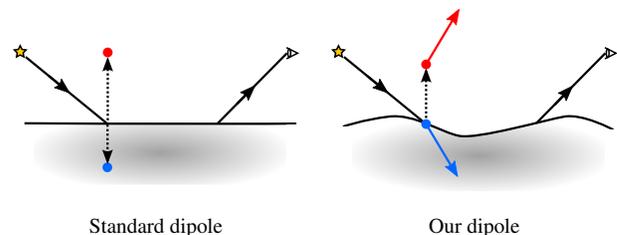


Figure 2: A standard dipole model uses two point sources to handle the boundary condition on a flat boundary. Our model uses two directional sources and relaxes the assumption of a flat boundary.

We introduce a directional subsurface scattering model which comes closer to path traced references. In particular, we relax the assumption on the direction of the incoming light. Our model is based on a new analytical solution to the diffusion equation for a ray of light in a highly scattering medium [Menon et al. 2005a; Menon et al. 2005b]. This solution is not directly applicable in computer graphics as it ignores boundary conditions. Inspired by the standard dipole model [Jensen et al. 2001], we extend this solution to a dipole construction, namely a mirrored pair of directional sources. We explain how to construct a practical BSSRDF based on this approach. Our final model is fully analytic, widely applicable to many different materials, and as computationally efficient as the standard dipole model with significantly improved accuracy. Figure 1 highlights our results.

2 Related Work

The most general approach for rendering participating media is to solve the radiative transfer equation by Monte Carlo ray tracing (path tracing) [Rushmeier 1988]. However, this approach becomes extremely costly for highly scattering media, where the average number of scattering events (and thus the number of rays to be traced) is very large. Hanrahan and Krueger [1993] pioneered rendering of highly scattering translucent materials in computer graph-

ics. They introduced analytical models for single scattering, but used path tracing for multiple scattering. Stam [1995] introduced diffusion theory to handle multiple scattering efficiently. He used a finite element method to solve the diffusion equation. A fully analytical model for subsurface scattering was first introduced to graphics by Jensen et al. [2001]. We refer to their model as the standard dipole in this paper.

There are a number of practical but approximate analytical solutions available for the diffusion equation. Jensen et al. [2001] used a solution for a point light in a semi-infinite medium. D'Eon and Irving [2011] and d'Eon [2012] recently presented more accurate models based on the same solution. The key difference in our model is that we use an approximate solution for a ray of light in an infinite medium [Menon et al. 2005a; Menon et al. 2005b] (see Figure 2). This solution enables us, for the first time in an analytical BSSRDF model, to take the directions of incident light rays into account.

Special case solutions for the diffusion equation are often derived for fluence inside a medium with no boundary. It is however important to consider boundary conditions in computer graphics, where we are interested in the radiance emerging at the surface. One popular approach to handle boundary conditions is the dipole approximation [Farrell et al. 1992; Jensen et al. 2001]. The dipole approximation was developed for handling light normally incident on a semi-infinite medium with a planar surface. The dipole model has been improved in a number of ways such as multipole [Wang 1998; Donner and Jensen 2005] and quadpole constructions [Kienle 2005; Donner and Jensen 2007]. Our model extends the above mentioned ray source solution to a dipole construction with a positive and a negative ray source (Figure 2).

Another class of methods in computer graphics uses precomputation. Donner et al. [2009] introduced a method which precomputes solutions to the radiative transfer equation as a canonical tabulated solution. Yan et al. [2012] proposed another precomputation method for highly scattering media using a spherical Gaussian approximation. Both approaches relax the assumptions made in the models based on diffusion theory at the cost of storing precomputed tables. Unlike this previous work, our model remains fully analytical without any precomputation.

3 Theory

We include subsurface scattering in a traditional rendering by using the general form of the rendering equation [Jensen et al. 2001]:

$$L_o(\mathbf{x}_o, \vec{\omega}_o) = L_e(\mathbf{x}_o, \vec{\omega}_o) + L_r(\mathbf{x}_o, \vec{\omega}_o) = L_e(\mathbf{x}_o, \vec{\omega}_o) + \int_A \int_{2\pi} S(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o, \vec{\omega}_o) L_i(\mathbf{x}_i, \vec{\omega}_i) (\vec{\omega}_i \cdot \vec{n}_i) d\omega_i dA, \quad (1)$$

where $L_o(\mathbf{x}_o, \vec{\omega}_o)$ is the outgoing (or emergent) radiance in the direction $\vec{\omega}_o$ from the location \mathbf{x}_o on the surface A of a medium, $L_i(\mathbf{x}_i, \vec{\omega}_i)$ is the radiance incident on the surface A at the location \mathbf{x}_i from the direction $\vec{\omega}_i$, and \vec{n}_i is the surface normal at \mathbf{x}_i . We solve this equation given L_e as an emission term, while L_r is the integral part which is referred to as the reflectance term. The function S is called a *Bidirectional Scattering-Surface Reflectance Distribution Function* (BSSRDF).

The BSSRDF is defined by the factor of proportionality between an element of emergent radiance $dL_r(\mathbf{x}_o, \vec{\omega}_o)$ and an element of incident flux $d\Phi_i(\mathbf{x}_i, \vec{\omega}_i)$ [Nicodemus et al. 1977]

$$S(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o, \vec{\omega}_o) = \frac{dL_r(\mathbf{x}_o, \vec{\omega}_o)}{d\Phi_i(\mathbf{x}_i, \vec{\omega}_i)}. \quad (2)$$

It is common to split the BSSRDF into a single scattering term $S^{(1)}$ and a multiple scattering term S_d [Hanrahan and Krueger 1993;



Figure 3: Path traced single scattering result (left) that should be added to the first three images in Figure 1 to get the full BSSRDF result. We also show a full result for the standard dipole as an example (right). Note that highlights on the back and the head are still significantly different from the path traced result.

Jensen et al. 2001], such that

$$S = T_{12}(S_d + S^{(1)})T_{21}, \quad (3)$$

where T_{12} and T_{21} are the Fresnel transmittance terms at the locations where the radiance enters and exits the medium, respectively.

There are a number of fast, approximate methods available for evaluating the single scattering term. In forward scattering materials, single scattering becomes increasingly unimportant as the asymmetry (g) increases, and most natural translucent materials are highly forward scattering [Frisvad et al. 2007]. The white grapefruit juice used in Figure 1 is also forward scattering, but with fairly low asymmetry ($g \approx 0.55$). Even in such a material, the contribution from single scattering is still low (see Figure 3). We therefore focus on the multiple scattering term for highly scattering media.

3.1 Diffusion Theory for a Ray of Light

Diffusion theory provides a way to derive analytic expressions for the multiple scattering term of the BSSRDF (S_d). This is usually done by finding a special case solution for the classic diffusion equation, which is [Ishimaru 1978]

$$(D\nabla^2 - \sigma_a)\phi_d(\mathbf{x}) = -q(\mathbf{x}) + 3D\nabla \cdot \mathbf{Q}(\mathbf{x}), \quad (4)$$

where $\phi_d(\mathbf{x}) = \int_{4\pi} L_d(\mathbf{x}, \vec{\omega}') d\omega'$ is the diffusive part of the fluence, σ_a is the absorption coefficient, D is the diffusion coefficient, and q and \mathbf{Q} are zeroth and first order source terms.

The standard dipole model [Jensen et al. 2001] uses the solution of the diffusion equation for a point source in an infinite medium. Consequently, the multiple scattering part of this model depends only on the distance between the points of incidence and emergence. The direction of the incident light ray affects only the Fresnel transmittance term. Our key contribution is a new BSSRDF model based on a more recent solution to the diffusion equation for a ray of light in an infinite medium [Menon et al. 2005a; Menon et al. 2005b]. Our model thus takes the directions of incident light rays into account.

Consider a ray which starts at the origin of the coordinate system and proceeds along the z -axis. The reduced intensity (or directly transmitted) light due to this source is

$$L_{ri}(\mathbf{x}, \vec{\omega}) = \Phi\delta(x)\delta(y)\Theta(z)e^{-\sigma_t z}\delta(\vec{\omega} - \vec{\omega}_z), \quad (5)$$

where Φ is the radiant flux of the source and σ_t is the extinction coefficient of the medium. We have set $\mathbf{x} = (x, y, z)$ while $\vec{\omega}_z$ is

the direction of the z -axis and $\Theta(z)$ is the Heaviside step function which is 1 when $z \geq 0$ and 0 otherwise. This equation leads to the following zeroth and first order source terms:

$$q(\mathbf{x}) = \sigma_s \Phi \delta(x) \delta(y) \Theta(z) e^{-\sigma_t z} \quad (6)$$

$$\mathbf{Q}(\mathbf{x}) = gq(\mathbf{x}) \vec{\omega}_z, \quad (7)$$

where σ_s is the scattering coefficient which is related to extinction and absorption by $\sigma_t = \sigma_s + \sigma_a$. The asymmetry parameter g is the mean cosine of the scattering angle.

Inserting the source terms (6–7) in the diffusion equation (4) and making the assumptions that lead to the usual dipole approximation, we reach an approximate directional solution for a ray of light in an infinite medium [Menon et al. 2005a; Menon et al. 2005b]:

$$\phi'_d(r, \theta) = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr} r}}{r} \left(1 + 3D \frac{1 + \sigma_{tr} r}{r} \cos \theta \right). \quad (8)$$

where $\sigma_{tr} = \sqrt{\sigma_a/D}$ is the effective transport coefficient, $r = |\mathbf{x}|$ is the distance from the point of incidence, and θ is the angle with the ray direction (see Figure 4):

$$\cos \theta = z/r = (\mathbf{x} \cdot \vec{\omega}_z)/r. \quad (9)$$

It is interesting to note that, for a ray of light which is perpendicular to the surface ($\cos \theta = 0$), Equation 8 turns into the solution for a point source which is used in existing analytical BSSRDF models [Jensen et al. 2001; d'Eon and Irving 2011]. Note that the assumptions are that we go to the asymptotic regions of the medium, far away from sources and boundaries where $r \gg 1/\sigma_s$, and we assume weak absorption $\sigma_a \ll \sigma_s$. The same assumptions have been employed in other analytical BSSRDF models. We use a prime ' to indicate that a solution is for an infinite medium (ϕ'_d above and also S'_d in the following).

Suppose we have incident radiance $L_i(\mathbf{x}_i, \vec{\omega}_i)$ and let it refract into a scattering material to get

$$L_t(\mathbf{x}_i, \vec{\omega}_{12}) = \eta^{-2} T_{12} L_i(\mathbf{x}_i, \vec{\omega}_i), \quad (10)$$

where $\eta = \eta_2/\eta_1$ is the relative refractive index, $\vec{\omega}_{12}$ is the direction of the transmitted ray given by the law of refraction, and T_{12} is the Fresnel transmittance. In an arbitrary coordinate system, with $\mathbf{x} = \mathbf{x}_o - \mathbf{x}_i$, we then have the following setup for Equation 8 (see Figure 4):

$$\Phi = T_{12} \Phi_i, \quad \vec{\omega}_z = \vec{\omega}_{12}, \quad r = |\mathbf{x}|, \quad \cos \theta = \frac{\mathbf{x} \cdot \vec{\omega}_{12}}{r}.$$

3.2 Emergent Radiance

To get an expression for the diffusive part of the BSSRDF, we must relate the fluence to the emergent radiance. The emergent radiance L_r due to diffuse subsurface scattering is given by

$$L_r(\mathbf{x}_o, \vec{\omega}_o) = \eta^2 T_{21} L_d(\mathbf{x}_o, \vec{\omega}_{21}), \quad (11)$$

where $\vec{\omega}_{21}$ is the direction of the ray from inside the medium refracting to the direction $\vec{\omega}_o$ according to the law of refraction and T_{21} is the Fresnel transmittance. Combining the diffusion approximation [Ishimaru 1978] with Fick's law of diffusion [Fick 1855], the diffusely scattered radiance is

$$L_d(\mathbf{x}, \vec{\omega}) = \frac{1}{4\pi} \phi_d(\mathbf{x}) - \frac{3}{4\pi} D(\mathbf{x}) \vec{\omega} \cdot \nabla \phi_d(\mathbf{x}) \quad (12)$$

with $D = 1/(3\sigma'_t)$, where $\sigma'_t = \sigma_s(1 - g) + \sigma_a$ is called the reduced extinction coefficient.

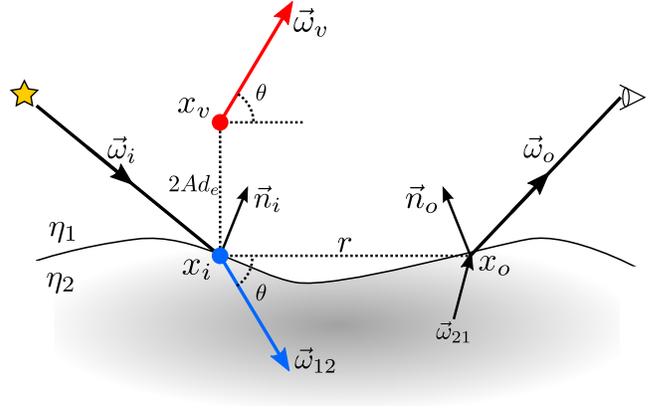


Figure 4: The dipole configuration of our model. It is similar to the dipole with point sources, but we have ray sources. The direction $\vec{\omega}_{12}$ of the refracted ray is used for the real source (blue). We mirror this direction in a modified tangent plane to find the direction $\vec{\omega}_v$ of the virtual source (red). The modified tangent plane contains $\mathbf{x}_o - \mathbf{x}_i$, and it is perpendicular to the plane spanned by \vec{n}_i and $\mathbf{x}_o - \mathbf{x}_i$. The origin of the virtual source is displaced along the normal of this modified plane.

The relation between L_d and the diffusive part of the BSSRDF S_d is given by Equations 2–3 and 11:

$$T_{12} S_d T_{21} = \eta^2 \frac{d(T_{21} L_d)}{d\Phi_i}. \quad (13)$$

Assuming that the diffusive light at the point of emergence no longer depends on the outgoing direction due to a large number of scatterings, we have $S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o, \vec{\omega}_o) = S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o)$. Integrating over outgoing directions, we get

$$T_{12} S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o) 4\pi C_\phi(1/\eta) = \frac{dM_d(\mathbf{x}_o)}{d\Phi_i(\mathbf{x}_i, \vec{\omega}_i)}, \quad (14)$$

where M_d is the diffuse radiant exitance and C_ϕ is a function of η which relates to hemispherical integration of the Fresnel transmittance. The η^2 factor disappears from the equation because we use $1/\eta$ as the argument for C_ϕ [Aronson 1995]. The diffuse radiant exitance is

$$M_d(\mathbf{x}_o) = \int_{2\pi} T_{21} L_d(\mathbf{x}_o, \vec{\omega}_{21}) (\vec{n}_o \cdot \vec{\omega}_o) d\omega_o, \quad (15)$$

where \vec{n}_o is the surface normal pointing outward at the point of emergence such that the integral is over the hemisphere with $\vec{n}_o \cdot \vec{\omega}_o > 0$. Inserting the expression for L_d (12), we get an integral that has been investigated many times before [Haskell et al. 1994; Aronson 1995; Kienle and Patterson 1997]. The solution is

$$M_d(\mathbf{x}_o) = C_\phi(\eta) \phi_d(\mathbf{x}_o) - C_E(\eta) D(\mathbf{x}_o) \vec{n}_o \cdot \nabla \phi_d(\mathbf{x}_o), \quad (16)$$

where C_ϕ and C_E are functions that d'Eon and Irving [2011] provide convenient polynomial approximations for. Equation 16 was introduced to graphics in the improved BSSRDF model by d'Eon and Irving [2011].

As opposed to previous work, we now use the diffuse fluence from

the directional solution (8). The gradient of this expression is

$$\nabla\phi'_d = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr}r}}{r^3} \left(\vec{\omega}_z 3D(1 + \sigma_{tr}r) - \mathbf{x}(1 + \sigma_{tr}r) - \mathbf{x} 3D \frac{3(1 + \sigma_{tr}r) + (\sigma_{tr}r)^2}{r} \cos\theta \right).$$

Inserting this solution and its gradient into Equation 16, we get an expression for the diffuse radiant exitance M_d . When this is inserted in Equation 14, we obtain an expression for the diffusive part of the BSSRDF inside an infinite medium:

$$\begin{aligned} S'_d(\mathbf{x}_i, \vec{\omega}_{12}; \mathbf{x}_o) 4C_\phi(1/\eta) = & \\ \frac{1}{4\pi^2} \frac{e^{-\sigma_{tr}r}}{r^3} \left[C_\phi(\eta) \left(\frac{r^2}{D} + 3(1 + \sigma_{tr}r)r \cos\theta \right) \right. & \\ - C_E(\eta) \left((\vec{\omega}_{12} \cdot \vec{n}_o) 3D(1 + \sigma_{tr}r) - (\mathbf{x}_o - \mathbf{x}_i) \right. & \\ \cdot \vec{n}_o \left. \left. \left((1 + \sigma_{tr}r) + 3D \frac{3(1 + \sigma_{tr}r) + (\sigma_{tr}r)^2}{r} \cos\theta \right) \right) \right], & \end{aligned} \quad (17)$$

where T_{12} and the flux disappear since $\Phi = T_{12}\Phi_i$, and we take the derivative with respect to Φ_i . The factor $4C_\phi(1/\eta)$ is the normalization factor also used by d'Eon and Irving [2011].

3.3 Boundary Conditions

The BSSRDF derived above (17) assumes that we have a ray of light in an *infinite medium*. Since we are interested in rendering media with boundaries, we need to take boundaries into account.

Dipole configuration. There are many sensible ways to incorporate the boundary in diffusion theory, and one has to make a number of decisions in order to build a new model [Haskell et al. 1994]. A common approach is to let the fluence vanish at the boundary of the medium. This is achieved by introducing a virtual source which is the real source mirrored in the surface tangent plane. This configuration is commonly referred to as a *dipole*. To make the diffusion approximation match transport theory better near the boundary, the fluence should instead vanish at an extrapolated boundary. The distance to this boundary is called the extrapolation distance d_e . According to exact transport theory, the extrapolation distance for a weakly absorbing medium with a planar surface is [Glasstone and Edlund 1952; Ishimaru 1978]

$$d_e = 0.7104/\sigma'_s = 2.121D/\alpha', \quad (18)$$

where $\sigma'_s = \sigma_s(1 - g)$ is the reduced scattering coefficient and $\alpha' = \sigma'_s/\sigma'_t$ is the reduced scattering albedo.

If the boundary of the medium has a mismatch between the outside and the inside refractive indices (η_1 and η_2 , respectively, see Figure 4), we must multiply the extrapolation distance by a reflection parameter A [Groenhuis et al. 1983]. When the real source is mirrored in the extrapolated boundary, the displacement of the virtual source becomes $2Ad_e$. The reflection parameter A is related to hemispherical integrals over the Fresnel transmittance. Using the coefficients from d'Eon and Irving [2011], it is

$$A = \frac{1 - C_E(\eta)}{2C_\phi(\eta)}. \quad (19)$$

Unlike the standard dipole model [Jensen et al. 2001], we have a dipole of ray sources. This difference influences how the sources should be positioned in the dipole. Figure 4 illustrates the configuration we have chosen.

Zero displacement of the real source. In the standard dipole model, the real source is placed at the first scattering event straight below the surface. This captures scattering due to a normally incident light ray. Since our source corresponds more accurately to the actual light ray, we should not displace the real source. This choice also requires us to clamp the distance $r = |\mathbf{x}_o - \mathbf{x}_i|$ to a minimum distance. This clamping is not only for avoiding numerical errors, but also for avoiding the region where the assumption of uniform emergent radiance is invalid. Our clamping is done by

$$r = \begin{cases} \max(|\mathbf{x}_o - \mathbf{x}_i|, d_e) & \vec{\omega}_{12} \cdot \vec{\omega}_{21} < 0 \text{ (front-lit)} \\ \max(|\mathbf{x}_o - \mathbf{x}_i|, D) & \text{otherwise (back-lit)} \end{cases}. \quad (20)$$

Mirroring of the direction. Since our source is directional, we also need to mirror the direction as well as the origin. While one obvious choice of mirror plane is the tangent plane defined by \vec{n}_i , we found that this choice leads to an unstable solution when a surface is curved or has small details. Instead, we use a modified tangent plane with normal

$$\vec{n}_i^* = \frac{\mathbf{x}_o - \mathbf{x}_i}{|\mathbf{x}_o - \mathbf{x}_i|} \times \frac{\vec{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)}{|\vec{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)|} \quad (21)$$

to mirror the source. If \mathbf{x}_o is in the surface tangent plane, this modified boundary equals the surface tangent plane.

Putting the above all together, our final BSSRDF is

$$S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o) = S'_d(\mathbf{x}_i, \vec{\omega}_{12}; \mathbf{x}_o) - S'_d(\mathbf{x}_v, \vec{\omega}_v; \mathbf{x}_o) \quad (22)$$

where, $\mathbf{x}_v = \mathbf{x}_i + 2Ad_e\vec{n}_i^*$ and $\vec{\omega}_v = \vec{\omega}_{12} - 2(\vec{\omega}_{12} \cdot \vec{n}_i^*)\vec{n}_i^*$ (see Figure 4).

4 Implementation

Just like the standard dipole model, our model can be implemented in a number of ways. The only major difference is that our model takes directions into account. Techniques which assume irradiance as input need minor modification in order to use our model. We implemented our model in the following two approaches.

Direct Monte Carlo integration. As described by Jensen et al. [2001], it is possible to use a BSSRDF in a Monte Carlo ray tracer. We integrate the BSSRDF in a progressive path tracer by distributing points evenly across the surface of the translucent object using a dart throwing technique. A new set of points is sampled iteratively. For every sampled surface point, the incident illumination is sampled from one direction. When a ray hits a translucent material, we loop over all surface samples and accept or reject them using a Russian roulette with the exponential term in the BSSRDF, $\exp(-\sigma_{tr}r)$, as the probability of acceptance.

Hierarchical integration. The hierarchical integration method for the standard dipole model [Jensen and Buhler 2002] works well with our model. Our implementation is almost the same as the original method. The only difference is that the irradiance computation is no longer separable from the evaluation of the BSSRDF. Each irradiance sample will be a list of differential irradiance samples instead of its sum at the same location. We also use the same list of directions across all irradiance samples in order to spatially cluster them without any change in the algorithm. Each evaluation of the BSSRDF now goes over the list and uses directions. The number of evaluations of the BSSRDF increases with this approach. One possible optimization would be to extend clustering to take into account directions, so that we can also cluster directional samples. We chose a simpler approach of utilizing the existing implementation of the hierarchical integration method.

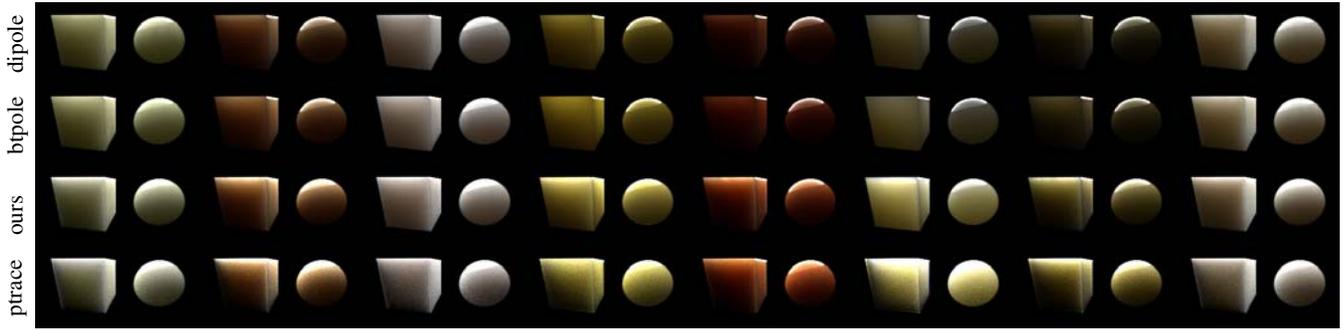


Figure 5: Simple test cases comparing the standard dipole model (*dipole*) [Jensen et al. 2001], the better dipole model (*btpole*) [d’Eon 2012], our model (*ours*), and the reference solutions rendered by path tracing (*ptrace*). Materials used are, from left to right, apple, chocolate milk (regular), marble, potato, skin1, soy milk (regular), white grapefruit juice, and whole milk all from corresponding measured values [Jensen et al. 2001; Narasimhan et al. 2006]. Results of our model are significantly closer to the path traced references while the rendering time, in most cases, only increases slightly.

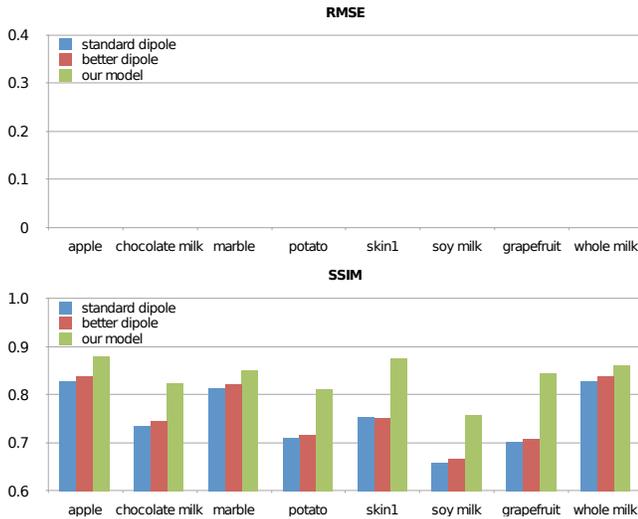


Figure 6: Objective image quality measurements for the renderings in Figure 5. Our model consistently outperforms other models.

5 Results

We implemented the standard dipole [Jensen et al. 2001], quantized diffusion [d’Eon and Irving 2011], the better dipole [d’Eon 2012], and our model all with the two approaches mentioned in the previous section. For material parameters, we use measured values of actual materials [Jensen et al. 2001; Narasimhan et al. 2006]. We exclude single scattering in all the renderings with our model as well as existing models since we only deal with highly scattering materials, and since it is expensive to render it with path tracing (16 hours for Figure 3, left). Note that adding single scattering does not recover the missing translucency effects in the existing models (shown in Figure 3), and our results can be similarly improved by adding single scattering.

To validate our model, we compare against full solutions of the radiative transfer equation using unbiased path tracing as described by Rushmeier [1988]. Figure 5 compares our model to the standard dipole and the better dipole for simple shapes of varying materials. To provide an objective comparison of the images, we measure the root-mean-squared error (RMSE) and the structural similarity index (SSIM [Wang et al. 2004]) of the different images using the path traced reference. Figure 6 plots all the measurements. Our method consistently achieves lower error and larger similarity.

We also ran similar comparisons for more complex geometries. Figure 1 shows the bunny model rendered with the white grapefruit juice material. The result of our model is visually closer to the path traced reference at the same cost as the other models. The RMSE validates that our model is indeed more accurate. Figure 7 shows the same scene configuration with the marble material. For this material, existing models tend to work rather well, yet our model is still more accurate than the other models. Figure 8 shows renderings of a chicken material which has lower albedo (≈ 0.7). The difference between our model and existing models is even more prominent. In particular, our model captures bright scattered highlights which are completely missing in existing models. The results in Figures 1, 5, 7, and 8 were rendered using direct Monte Carlo integration. The remaining results were rendered using hierarchical integration.

Figure 10 is a comparison with quantized diffusion [d’Eon and Irving 2011]. It illustrates that, except for subtle differences, quantized diffusion generally give us results that are very similar to the results obtained with the better dipole [d’Eon 2012]. Our model on the other hand gives us significantly different results which we have confirmed to be more accurate in all comparisons to path tracing references. Figure 11 demonstrates renderings with image based lighting. All the images in this figure were rendered in less than 3 min on a regular desktop PC. However, renderings using the better dipole took less than 1 min, as we do not need to sample directions for every shading point. Note that our model still works well under complex illumination and across a wide range of materials.

Since our model has a more complicated mathematical expression, it requires some extra computations. In a simple scene where the ray tracing is cheap, evaluation of our model can be up to 1.5 times slower per sampled path. In addition, since our model captures more features, it sometimes requires more samples to converge. All together, the rendering time can sometimes be up to 2 times higher before a converged image is obtained.

6 Discussion

Full planar model and modified boundary plane. While our model uses distances and angles computed directly from given locations and directions, we found that we could also assume a completely planar configuration and compute distances and angles accordingly. For example, we can compute the distance to the virtual source differently. This configuration is employed in the standard dipole model [Jensen et al. 2001] and in quantized diffusion [d’Eon and Irving 2011]. Our early experiments however show that this model is slightly inferior to the current full 3D model (see Figure 9, center). We also experienced that using the modified bound-

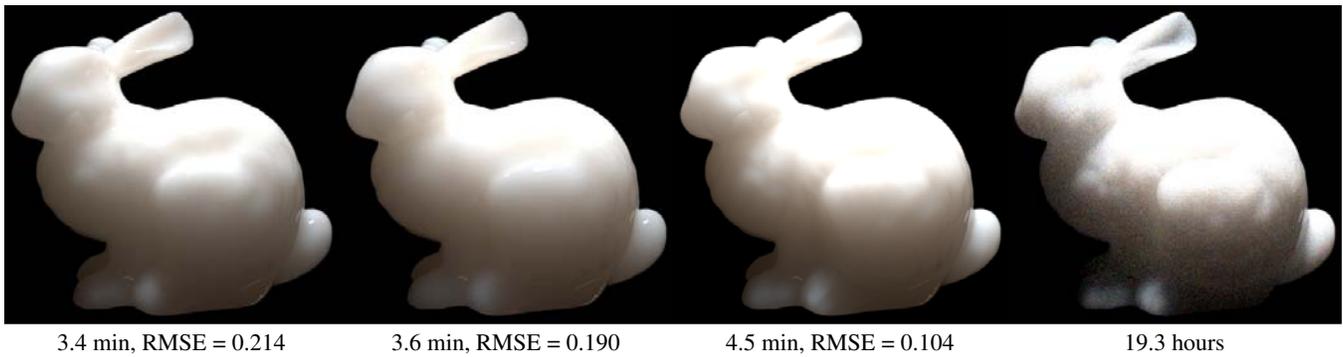


Figure 7: Same scene configuration as Figure 1 but with marble [Jensen et al. 2001]. Although existing models work rather well for this material, our model is still closer to the path traced reference solution.

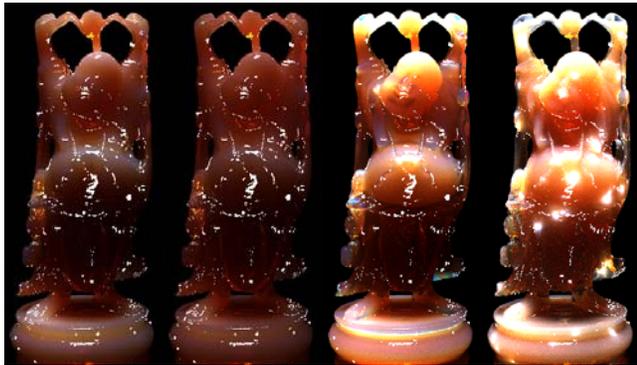


Figure 8: The buddha model made out of chicken meat [Jensen et al. 2001] and rendered using, from left to right, the standard dipole, the better dipole, our directional dipole, and path tracing. Our model is closer to the reference. The specular highlights have high intensity in this case and they cause volume caustics which appear only in path tracing.

ary plane to mirror sources is very important. Without this modification, our model does not work well (see Figure 9, right).

Reciprocity. Unlike the standard dipole model which only depends on the distance, our model is not precisely reciprocal (i.e., swapping variables for an incident point and an emergent point does not result in the same value). However, we have found that our model is very close to reciprocal in practice. We have seen little changes in rendered images when swapping variables. If we need a strictly reciprocal model, we can simply take the average of two evaluations of our model with swapped variables. All the results in this paper do not use this trick as the evaluation cost doubles.

Negative values. In rare cases, our BSSRDF (Equation 22) can return a negative value. It is a consequence of the approximations in the ray source solution (Equation 8). Our current solution is to simply clamp the final value to zero if it is negative. We have not found a case where clamping causes visible artifacts.

Limitations. Since we still make a number of assumptions to derive our model, it does not perfectly match the results of path tracing. For example, our model does not capture effects due to anisotropic scattering, as then the emergent radiance distribution is often not uniform. We also use a dipole to handle boundary conditions. This approach is known to have a number of failure cases

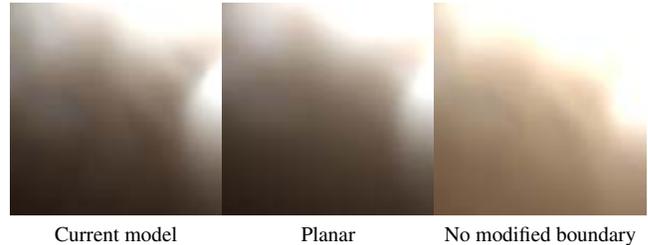


Figure 9: Comparison of our current model, the planar model, and our model without the modified boundary plane. The images are close ups of the bunny in Figure 7 without Fresnel reflections to highlight the differences. The planar model is fairly close to the current model, but missing subtle details. The use of the modified boundary plane is important for our model to work well.

such as sharp boundaries and thin features. Extensions of our model to quadpole and multipole constructions is interesting future work.

7 Conclusion

We introduced a novel BSSRDF model which comes closer to path traced references than existing models. Our model is built upon a new analytical solution to the diffusion equation which allows us to take the direction of incoming light into account. We explained how to derive a practical BSSRDF model by extending this solution into a dipole of real and virtual ray sources. Our work significantly improves the accuracy of an analytical BSSRDF model based on the diffusion theory. In particular, our model is capable of reproducing small details that are missing in existing models. We believe that our model is a valuable addition to the existing set of analytical BSSRDF models.

Acknowledgments

The work was in part supported by the Digital Prototypes project funded by the Danish Council for Technology and Innovation (Resultatkontrakt). Thanks to Jesper Mosegaard for encouragement and funding through the Digital Prototypes project. We would also like to thank Eugene d'Eon for helping us implement the quantized diffusion model.

References

- ARONSON, R. 1995. Boundary conditions for diffusion of light. *Journal of the Optical Society of America A* 12, 11 (November), 2532–2539.



Figure 10: Rendered images of the happy buddha with various materials. Each row is rendered with the better dipole [d'Eon 2012] (top), quantized diffusion [d'Eon and Irving 2011] (middle), and our model (bottom). The list of materials are the same as the simple test cases in Figure 5. Only our model results in very different images due to its directionality.

D'EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2011)* 30, 4 (July), 56:1–56:13.

D'EON, E. 2012. A better dipole (a publicly available manuscript). <http://www.eugenedeon.com/papers/betterdipole.pdf>.

DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2005)* 24, 3 (July), 1032–1039.

DONNER, C., AND JENSEN, H. W. 2007. Rendering translucent materials using photon diffusion. In *Proceedings of Eurographics Symposium on Rendering (Rendering Techniques '07)*, 243–251.

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

FARRELL, T. J., PATTERSON, M. S., AND WILSON, B. 1992. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue op-

tical properties *in vivo*. *Medical Physics* 19, 4 (July/August), 879–888.

FICK, A. 1855. On liquid diffusion. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science X*, 30–39. Abstracted by the author from the German original: Über Diffusion, Poggendorff's *Annalen der Physik und Chemie*, Vol. 94, pp. 59–86, 1855.

FRISVAD, J. R., CHRISTENSEN, N. J., AND JENSEN, H. W. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2007)* 26, 3 (July), 60:1–60:10.

GLASSTONE, S., AND EDLUND, M. C. 1952. *The Elements of Nuclear Reactor Theory*. D. van Nostrand Company, Inc., Princeton, New Jersey.

GROENHUIS, R. A. J., FERWERDA, H. A., AND TEN BOSCH, J. J. 1983. Scattering and absorption of turbid materials determined from reflection measurements. 1: Theory. *Applied Optics* 22, 16 (August), 2456–2462.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Computer Graphics (Proceedings of ACM SIGGRAPH 93)* (August), 165–174.



Figure 11: Image based lighting of the happy buddha via hierarchical integration. The images in the top row are rendered with the better dipole [d'Eon 2012] and the images in the bottom row are rendered with our model. The original dipole model [Jensen et al. 2001] and quantized diffusion [d'Eon and Irving 2011] gave almost the same results as the better dipole in this scene. The list of materials are the same as the simple test cases in Figure 5.

- HASKELL, R. C., SVAASAND, L. O., TSAY, T.-T., FENG, T.-C., MCADAMS, M. S., AND TROMBERG, B. J. 1994. Boundary conditions for the diffusion equation in radiative transfer. *Journal of the Optical Society of America A* 11, 10 (October), 2727–2741.
- ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media*. Academic Press, New York. Reissued by IEEE Press and Oxford University Press 1997.
- JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2002)* 21, 3 (July), 576–581.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, 511–518.
- KIENLE, A., AND PATTERSON, M. S. 1997. Improved solutions of the steady-state and the time-resolved diffusion equations for reflectance from a semi-infinite turbid medium. *Journal of the Optical Society of America A* 14, 1 (January), 246–254.
- KIENLE, A. 2005. Light diffusion through a turbid parallelepiped. *Journal of the Optical Society of America A* 22, 9 (September), 1883–1888.
- MENON, S., SU, Q., AND GROBE, R. 2005. Determination of g and μ using multiply scattered light in turbid media. *Physical Review Letters* 94 (April), 153904.
- MENON, S., SU, Q., AND GROBE, R. 2005. Generalized diffusion solution for light scattering from anisotropic sources. *Optics Letters* 30, 12 (June), 1542–1544.
- NARASIMHAN, S. G., GUPTA, M., DONNER, C., RAMAMOORTHY, R., NAYAR, S. K., AND JENSEN, H. W. 2006. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2006)* 25, 3 (July), 1003–1012.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometrical considerations and nomenclature for reflectance. Tech. rep., National Bureau of Standards (US), October.
- RUSHMEIER, H. E. 1988. *Realistic Image Synthesis for Scenes with Radiatively Participating Media*. PhD thesis, Cornell University.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques '95*, Springer, P. Hanrahan and W. Purgathofer, Eds., 41–50. Proceedings of the Sixth Eurographics Workshop on Rendering.
- WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (April), 600–612.
- WANG, L. V. 1998. Rapid modeling of diffuse reflectance of light in turbid slabs. *Journal of the Optical Society of America A* 15, 4 (April), 936–944.
- YAN, L.-Q., ZHOU, Y., XU, K., AND WANG, R. 2012. Accurate translucent material rendering under spherical gaussian lights. *Computer Graphics Forum* 31, 7 (September), 2267–2276.