

Directional Dipole Model for Subsurface Scattering

JEPPE REVALL FRISVAD

Technical University of Denmark

TOSHIYA HACHISUKA

Aarhus University

and

THOMAS KIM KJELDEN

The Alexandra Institute

Rendering translucent materials using Monte Carlo ray tracing is computationally expensive due to a large number of subsurface scattering events. Faster approaches are based on analytical models derived from diffusion theory. While such analytical models are efficient, they miss out on some translucency effects in the rendered result. We present an improved analytical model for subsurface scattering that captures translucency effects present in the reference solutions but remaining absent with existing 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 with a mirrored pair of directions. Our model is as computationally efficient as existing models while it includes single scattering without relying on a separate Monte Carlo simulation, and the rendered images are significantly closer to the references. Unlike some previous work, our model is fully analytic and requires no precomputation.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

General Terms: Theory, Algorithms

Additional Key Words and Phrases: Rendering, BSSRDF, diffusion dipole, radiative transfer, subsurface scattering, translucent materials, turbid media.

J. R. Frisvad and T. K. Kjeldsen acknowledge support from the Digital Prototypes project funded by the Danish Council for Technology and Innovation (Resultatkontrakt).

Authors' addresses: J. R. Frisvad (corresponding author), Technical University of Denmark, Richard Petersens Plads, Building 321, 2800 Kgs. Lyngby, Denmark; email: jerf@dtu.dk; T. Hachisuka Aarhus University, Nordre Ringgade 1, 8000 Aarhus C, Denmark; T. K. Kjeldsen, The Alexandra Institute, Aabogade 34, 8200 Aarhus N, Denmark.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

2014 Copyright held by the Owner/Author. Publication Rights Licensed to ACM. This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Transactions on Graphics 35, 1 (November 2014).

DOI: <http://dx.doi.org/10.1145/2682629>

ACM Reference Format:

Jeppe Revall Frisvad, Toshiya Hachisuka, and Thomas Kim Kjeldsen. 2014. Directional dipole model for subsurface scattering. *ACM Trans. Graph.* 34, 1, Article 5 (November 2014), 12 pages. DOI: <http://dx.doi.org/10.1145/2682629>

1. INTRODUCTION

Rendering of translucent materials, such as skin, foods, minerals, 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 a fast practical way of rendering such materials. This standard dipole model is, however, built upon a number of assumptions that 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 recently introduced analytical models [d'Eon and Irving 2011]. The lighting distribution in realistic scene configurations, however, is rarely directionally uniform.

We introduce a directional dipole model for subsurface scattering that produces images closer to path-traced references. Our model is a fully analytical BSSRDF that does not assume directionally uniform incident illumination. The derivation of our model is based on a new analytical solution to the diffusion equation for a ray of light in a highly scattering medium with infinite extent [Menon et al. 2005a; 2005b]. To handle boundary conditions, we extend this ray source solution to a dipole construction, namely a mirrored pair of directional sources. This construction is inspired by the standard dipole model [Jensen et al. 2001].

Conventionally, single scattering is separated out from a model for subsurface scattering. The ray source we use enables us to include approximate single scattering in our model. We can thereby evaluate complete subsurface scattering effects without relying on a separate, potentially costly evaluation of single scattering based on Monte Carlo ray tracing. Thus, while our inclusion of directional effects adds some costs in a rendering, the inclusion of single scattering balances these additional expenses, and outweighs them when light sources are small (or singular). 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]. This approach becomes extremely costly for highly scattering media, where the average number of

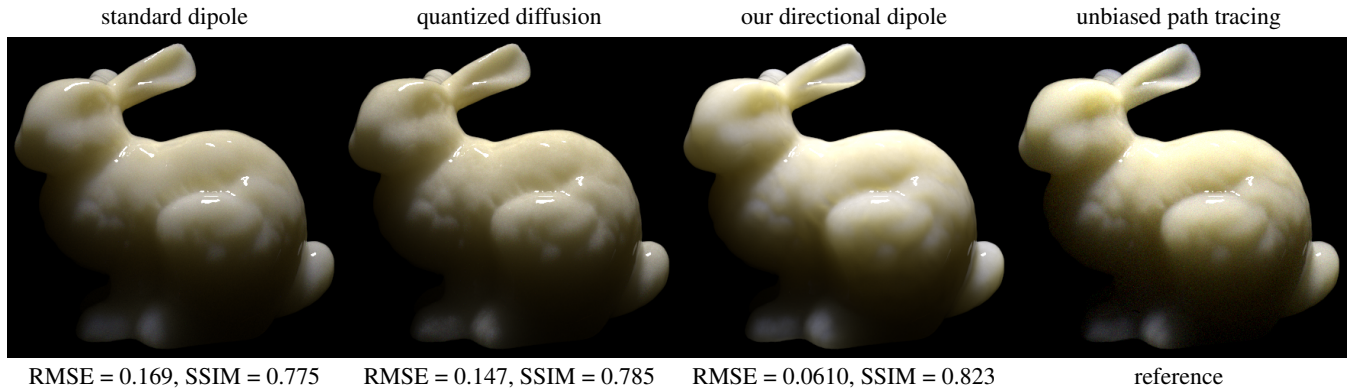


Fig. 1. Stanford Bunny made of white grapefruit juice [Narasimhan et al. 2006]. It was rendered using different models for subsurface scattering. In the standard dipole [Jensen et al. 2001] and in quantized diffusion [d’Eon and Irving 2011], accurate single scattering was added using unbiased path tracing [Rushmeier 1988]. Our model captures translucency effects that are present in the full path tracing but not in the results of the other analytical BSSRDF models. The Stanford Bunny model is courtesy of the Stanford University Computer Graphics Laboratory (<http://graphics.stanford.edu/data/3Dscanrep/>).

scattering events (and thus the number of rays to be traced) is very large. Hanrahan and Krueger [1993] pioneered efficient rendering of highly scattering translucent materials. They combined analytical single scattering in layers [Blinn 1982] with refractive boundaries and path-traced multiple scattering. Stam [1995] introduced diffusion theory to evaluate multiple scattering using a finite element method. Dorsey et al. [1999] adapted a general Monte Carlo method called volume photon mapping [Jensen and Christensen 1998] for subsurface scattering. A fully analytical BSSRDF for subsurface scattering was first introduced to graphics by Jensen et al. [2001]. We refer to their model as the standard dipole in this article. Our model is similar to the standard dipole in the sense that it is also fully analytic and does not use numerical approximation.

The standard dipole was combined with photon mapping by Donner and Jensen [2007]. They introduced an extended source model, which means that the standard dipole solution is integrated along the refracted ray. This method is referred to as photon diffusion. Although photon diffusion can take oblique incidence into account, the additional integration requires a larger number of samples than an analytical model such as the standard dipole or our model. D’Eon and Irving [2011] recently introduced an analytical extended source model. As in the standard dipole, they assume that light is normally incident. To solve the integral along the ray in the normal direction, they use a functional approximation based on a sum of Gaussians. In work concurrent with ours, Habel et al. [2013] showed that such integrals can be approximated efficiently using Monte Carlo integration instead of this functional approximation. Their algorithm is similar to photon diffusion and therefore takes oblique incidence into account. The key difference in our model is that we start from a fundamentally different solution of the diffusion equation for a ray of light in an infinite medium [Menon et al. 2005a; 2005b]. This solution enables us to derive the first analytical BSSRDF model that can take the directions of incident light rays into account. Unlike the other methods discussed here, our method does not require an expensive integration of the solution along the refracted ray.

Special case solutions for the diffusion equation are often derived for fluence (*scalar irradiance* in a volume) 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 boundary. One popular approach to handle boundary conditions is the dipole approximation [Farrell et al. 1992; Jensen et al. 2001]. The dipole approximation was originally developed for




Fig. 2. A standard dipole model uses two point sources to handle boundary conditions. The points are displaced along the normal at the point of incidence. Our model uses two directional sources that are displaced along the normal of a plane containing the points of incidence and emergence.

handling light normally incident on a semi-infinite medium with a planar surface. We relax the assumption of normal incidence.

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 a directional source solution [Menon et al. 2005a; 2005b] to a dipole construction with a positive and a negative ray source (Figure 2). To the best of our knowledge, our work introduces a dipole extension with directional sources for the first time.

While our focus in this article is to find an improved analytical model for subsurface scattering, another approach is to use precomputation. Donner et al. [2009] introduced a method that 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 these methods, our model remains fully analytic. As such, the implementation of our model is trivial and incurs no extra memory cost or additional precomputation time.

3. THEORY

To have subsurface scattering in a surface rendering method, we use 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 . The term L_e is the emitted radiance while L_r is the reflected radiance. The function S is 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]

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

Typically, the BSSRDF is split into three terms: reduced intensity (direct transmission), single scattering, and multiple scattering [Hanrahan and Krueger 1993; Jensen et al. 2001]. Previous models for subsurface scattering focus on the definition of the multiple scattering term and leave the remaining terms to be computed separately using other methods (or the other way around [Hanrahan and Krueger 1993]). However, this separation is not always convenient. To date, there is seemingly no accurate and efficient method for rendering single scattering in translucent materials (that is, optically thick media with smooth, refractive boundaries). We are only aware of approximate methods for single scattering that assume layers with $\mathbf{x}_o \approx \mathbf{x}_i$ [Blinn 1982; Hanrahan and Krueger 1993], approximate transmission distance [Jensen et al. 2001], nonrefractive boundaries [Sun et al. 2005], an optically thin medium [Walter et al. 2009], or rough boundaries [Habel et al. 2013]. To get accurate single scattering, we have to rely on path tracing, which can be computationally expensive.

Instead of splitting the BSSRDF into three terms, we formulate our model such that it includes single scattering based on the delta-Eddington (δE) approximation [Joseph et al. 1976]. In this approximation, the part of the single scattering that continues along the refracted ray is moved to the reduced intensity term $S_{\delta E}$. We then consider multiple scattering and the remaining single scattering to be the diffusive part of the BSSRDF S_d . The full BSSRDF in this approximation becomes

$$S = T_{12}(S_d + S_{\delta E})T_{21} , \quad (3)$$

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

The modified reduced intensity term $S_{\delta E}$ is defined as the usual reduced intensity term but using a modified set of scattering properties [Joseph et al. 1976]

$$\tilde{\sigma}_s = \sigma_s(1 - g^2) , \quad \tilde{\sigma}_t = \tilde{\sigma}_s + \sigma_a , \quad \tilde{g} = g/(g + 1) , \quad (4)$$

where σ_a is the absorption coefficient, σ_s is the scattering coefficient, and g is the asymmetry parameter (the mean cosine of the scattering angle). These δE scattering properties replace the corresponding scattering properties in the standard definition of the reduced intensity term. We return to the definition of this modified reduced intensity in Sections 3.1 and 3.2.

Overview. In the following, we describe how we derive S_d . The main idea is to consider both the origin and the direction of the refracted light ray within the medium. We can then use a solution to the diffusion equation where this ray is the source term. This solution gives us a formula for estimating the fluence at any point in the surrounding medium. Unlike existing models that use point sources, this formula is not only a function of distance, but of the angle between the ray direction and the direction toward the position that we

are interested in also. Once this solution is established, we consider boundary conditions by introducing a dipole with ray sources, and use the formula to find the radiance emerging from the medium at a point on the boundary. This leads to a new analytical BSSRDF that takes the direction of the incident light and the direction toward the point of emergence into account.

3.1 Diffusion Theory for a Ray of Light

As in previous work, we derive the diffusive part of the BSSRDF S_d using diffusion theory. This is usually done by finding a special case solution for the classic diffusion equation, which is [Ishimaru 1978]

$$(D\nabla^2 - \sigma_a)\phi_d(\mathbf{x}) = -q(\mathbf{x}) + 3D\nabla \cdot \mathbf{Q}(\mathbf{x}) , \quad (5)$$

where $\phi_d(\mathbf{x}) = \int_{4\pi} L_d(\mathbf{x}, \vec{\omega}') d\omega'$ is the diffusive part of the fluence, 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 diffusive 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; 2005b]. Our model thus takes the directions of incident light rays into account.

Consider a ray from the origin along the z -axis in a Cartesian xyz coordinate system. In the δE approximation, the reduced intensity (or directly transmitted) light due to this source is

$$L_{ri}(\mathbf{x}, \vec{\omega}) = \Phi\delta(x)\delta(y)\Theta(z)e^{-\tilde{\sigma}_t z}\delta(\vec{\omega} - \vec{\omega}_z) , \quad (6)$$

where Φ is the radiant flux of the source. 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 (it defines a half-line starting at the origin). This ray source leads to the following zeroth- and first-order source terms.

$$q(\mathbf{x}) = \tilde{\sigma}_s\Phi\delta(x)\delta(y)\Theta(z)e^{-\tilde{\sigma}_t z} \quad (7)$$

$$\mathbf{Q}(\mathbf{x}) = \tilde{g}q(\mathbf{x})\vec{\omega}_z . \quad (8)$$

These terms would be simpler if we were using an isotropic point source, especially the first-order source term $\mathbf{Q}(\mathbf{x})$, which would then be zero. With Eq. (8), we can more accurately represent the anisotropy of the radiance. The region around the origin where anisotropy is important grows with the asymmetry parameter g [Bevilacqua and Depeursing 1999]. This anisotropy is not accurately represented by the extended source models [Donner and Jensen 2007; d'Eon and Irving 2011; Habel et al. 2013] since they rely on (an integration of) the isotropic point source solution.

As in the standard dipole, we assume that the point of estimation is not too close to sources and boundaries ($|\mathbf{x}| \gg 1/\sigma_s$) and that the medium is a weak absorber ($\sigma_a \ll \sigma_s$). Inserting the source terms (7 and 8) into the diffusion equation (5) and making these assumptions, we reach an approximate directional solution for a ray of light in an infinite medium [Menon et al. 2005a; 2005b]

$$\phi'_d(r, \theta) = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr} r}}{r} \left(1 + 3D \frac{1 + \sigma_{tr} r}{r} \cos \theta \right) , \quad (9)$$

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 3);

$$\cos \theta = z/r = (\mathbf{x} \cdot \vec{\omega}_z)/r . \quad (10)$$




Fig. 3. 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 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. Note that $\theta \neq \theta_v$ and $r \neq r_v$.

It is interesting to note that, for a ray of light perpendicular to the direction toward the point of interest ($\cos \theta = 0$), Eq. (9) turns into the point source solution which is used in existing analytical BSSRDF models [Jensen et al. 2001; d'Eon and Irving 2011]. We use a prime ' to indicate that a solution disregards boundaries (ϕ'_d before 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 the transmitted radiance

$$L_t(\mathbf{x}_i, \vec{\omega}_{12}) = \eta^{-2} T_{12} L_i(\mathbf{x}_i, \vec{\omega}_i), \quad (11)$$

where $\eta = \eta_2/\eta_1$ is the relative refractive index and $\vec{\omega}_{12}$ is the direction of the refracted ray given by the law of refraction applied to $\vec{\omega}_i$. To use a refracted ray of arbitrary origin and direction as the ray source in our model, we have $\mathbf{x} = \mathbf{x}_o - \mathbf{x}_i$ and the following setup for Eq. (9) (see Figure 3).

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

3.2 Emergent Radiance

We now briefly return to the reduced intensity term and then describe how we obtain the emergent radiance from the fluence using approximations that were also used by d'Eon and Irving [2011].

Due to the delta-Eddington approximation, the transmitted radiance (11) and the radiance due to δE reduced intensity (6) enable us to indirectly define the modified reduced intensity part of the BSSRDF $S_{\delta E}$. The δE modified direct transmission emerging from the medium is

$$L_{r,\delta E}(\mathbf{x}_o, \vec{\omega}_o) = \eta^2 T_{21} e^{-\tilde{\sigma}_t s} L_t(\mathbf{x}_{ri}, \vec{\omega}_{21}), \quad (12)$$

where $\mathbf{x}_{ri} = \mathbf{x}_o - s \vec{\omega}_{21}$ is the only point on the surface that can contribute to this reduced intensity term, s is the distance from \mathbf{x}_o to this point, and $\vec{\omega}_{21}$ is the direction of the ray from inside the medium refracting to the direction $\vec{\omega}_o$ according to the law of refraction. This term is easily evaluated using traditional ray tracing.

To get an expression for the diffusive part of the BSSRDF, we must relate the fluence to the emergent radiance. The emergent

radiance $L_{r,d}$ due to diffuse subsurface scattering is given by

$$L_{r,d}(\mathbf{x}_o, \vec{\omega}_o) = \eta^2 T_{21} L_d(\mathbf{x}_o, \vec{\omega}_{21}). \quad (13)$$

Combining the diffusion approximation [Ishimaru 1978] with Fick's law of diffusion [Fick 1855], the diffusely scattered radiance is

$$L_d(\mathbf{x}, \vec{\omega}) = \frac{1}{4\pi} \phi_d(\mathbf{x}) - \frac{3}{4\pi} D \vec{\omega} \cdot \nabla \phi_d(\mathbf{x}) \quad (14)$$

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

The relationship between L_d and the diffusive part of the BSSRDF S_d is given by Eqs. (2), (3), and (13):

$$T_{12} S_d T_{21} = \eta^2 \frac{d(T_{21} L_d)}{d\Phi_i}. \quad (15)$$

Assuming that the diffusive light at the point of emergence no longer depends on the outgoing direction due to a large number of scattering events, 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

$$T_{12} S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o) 4\pi C_\phi(1/\eta) = \frac{dM_d(\mathbf{x}_o)}{d\Phi_i(\mathbf{x}_i, \vec{\omega}_i)}, \quad (16)$$

where M_d is the diffuse radiant exitance and C_ϕ is a function of η that relates to hemispherical integration of the Fresnel transmittance (see Appendix A). The η^2 factor cancels out because $1/\eta$ was chosen as the argument of C_ϕ [Aronson 1995]. The diffuse radiant exitance is

$$M_d(\mathbf{x}_o) = \int_{2\pi} T_{21} L_d(\mathbf{x}_o, \vec{\omega}_{21}) (\vec{n}_o \cdot \vec{\omega}_o) d\omega_o, \quad (17)$$

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 (14), we get an integral that has been investigated many times before [Haskell et al. 1994; Aronson 1995; Kienle and Patterson 1997]. The solution is

$$M_d(\mathbf{x}_o) = C_\phi(\eta) \phi_d(\mathbf{x}_o) - C_E(\eta) D \vec{n}_o \cdot \nabla \phi_d(\mathbf{x}_o), \quad (18)$$

where C_ϕ and C_E are in Appendix A. Eq. (18) was also used in the BSSRDF model of d'Eon and Irving [2011].

As opposed to previous work, we now use the diffuse fluence from the directional solution (9). The gradient of this expression is

$$\nabla \phi'_d = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr} r}}{r^3} \left(\vec{\omega}_{12} 3D(1 + \sigma_{tr} r) - \mathbf{x}(1 + \sigma_{tr} r) - \mathbf{x} 3D \frac{3(1 + \sigma_{tr} r) + (\sigma_{tr} r)^2}{r} \cos \theta \right). \quad (19)$$

Inserting the directional solution (9) and its gradient (19) into Eq. (18), we get an expression for the diffuse radiant exitance M_d . When this is inserted in Equation 16, we obtain an expression for the diffusive part of the BSSRDF when disregarding boundaries

$$S'_d(\mathbf{x}, \vec{\omega}_{12}, r) = \frac{1}{4C_\phi(1/\eta)} \frac{1}{4\pi^2} \frac{e^{-\sigma_{tr} r}}{r^3} \left[C_\phi(\eta) \left(\frac{r^2}{D} + 3(1 + \sigma_{tr} r) \mathbf{x} \cdot \vec{\omega}_{12} \right) - C_E(\eta) \left(3D(1 + \sigma_{tr} r) \vec{\omega}_{12} \cdot \vec{n}_o - \left((1 + \sigma_{tr} r) + 3D \frac{3(1 + \sigma_{tr} r) + (\sigma_{tr} r)^2}{r^2} \mathbf{x} \cdot \vec{\omega}_{12} \right) \mathbf{x} \cdot \vec{n}_o \right) \right], \quad (20)$$

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

3.3 Boundary Conditions

The BSSRDF derived earlier (20) disregards boundaries in the sense that it was derived for a ray of light in an *infinite medium*. Since our purpose is to render media with boundaries, we need to take boundary conditions into account. Like existing BSSRDF models [Jensen et al. 2001; d'Eon and Irving 2011], we assume a semi-infinite medium with a flat boundary and apply the resulting model to arbitrary boundary conditions.

We note that any approximate boundary conditions, including ours and the ones used for existing models, are more or less empirically based. It is widely acknowledged [Aronson 1995, for example] that comparison with full solutions of the radiative transfer equation is the only reasonable way to evaluate the appropriateness of approximate boundary conditions. In the following, we provide reasoning behind each of our choices but only as motivation. We validate our final model via numerical experiments in Section 5.

3.3.1 Dipole Configuration. There are many sensible ways to incorporate the boundary in diffusion theory, and one has to make a number of decisions when building a new model [Haskell et al. 1994]. One common approach is to let the fluence vanish at the boundary of the medium. This is achieved by introducing a virtual source that is the real source mirrored in the surface tangent plane, a configuration commonly referred to as a *dipole*. Unlike the standard dipole [Jensen et al. 2001], we have a dipole of ray sources which influences how the sources should be positioned in the dipole. Figure 3 illustrates the configuration we have chosen.

The diffusion approximation matches the exact solution better if the fluence vanishes at an *extrapolated boundary*, that is, a boundary which is the same as the actual boundary but displaced outward in the normal direction [Glasstone and Edlund 1952; Davison 1958; Haskell et al. 1994; Aronson 1995]. The orthogonal distance from the actual boundary to this extrapolated boundary is called the extrapolation distance d_e . Existing models use $d_e = 2D$. This standard approximation for d_e was originally derived from diffusion theory for nonabsorbing media with nonrefractive boundaries ($\eta = 1$). Comparing with exact calculations of d_e for this and other types of media, Aronson [1995] found that $2D$ is imprecise in all cases. We thus use an approximation that is closer to the right d_e for weakly absorbing media with non-refractive boundary [Davison 1958]

$$d_e = 2.131 D / \sqrt{\alpha'} , \quad (21)$$

where $\alpha' = \sigma'_s / \sigma'_t$ is the reduced scattering albedo, which is the ratio of the reduced scattering coefficient $\sigma'_s = \sigma_s(1 - g)$ to the reduced extinction coefficient $\sigma'_t = \sigma'_s + \sigma_a$.

If the boundary of the medium has a mismatch of refractive indices (η_1 and η_2 ; see Figure 3), 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 [Aronson 1995]. We use the same reflection parameter as d'Eon and Irving [2011]. A convenient way to express A is

$$A(\eta) = \frac{1 - C_E(\eta)}{2C_\phi(\eta)} . \quad (22)$$

3.3.2 Modified Tangent Plane. Since our source is directional, we also need to mirror the direction as well as the origin. While

one obvious choice of mirroring plane is the tangent plane defined by \vec{n}_i , we found this choice leads to an inaccurate solution if the assumption of a flat boundary is violated (which in practice is almost always the case). Instead, we mirror the source using a modified tangent plane with normal

$$\vec{n}_i^* = \begin{cases} \vec{n}_i & , \text{ for } \mathbf{x}_o = \mathbf{x}_i \\ \frac{\mathbf{x}_o - \mathbf{x}_i}{|\mathbf{x}_o - \mathbf{x}_i|} \times \frac{\vec{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)}{|\vec{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)|} & , \text{ otherwise } \end{cases} . \quad (23)$$

As we will demonstrate in Section 6.2, this choice in particular eliminates issues in backlit regions. These issues are due to exaggerated directional effects caused by the assumption of a semi-infinite medium. The modification basically ensures that $\cos \theta$ only changes in the virtual source as if the geometry were planar. Note that the modified normal is used only to set the position and the direction of the virtual source. The normal used for Fresnel and cosine terms is not affected. If the boundary is actually flat, or, more generally, if the point of emergence lies in the tangent plane defined by \vec{n}_i , the modified normal equals the normal at the point of incidence.

3.3.3 Distance to the Real Source. In the standard dipole, the real source is placed at the average depth $z_r = 1/\sigma'_t = 3D$ of the first scattering event straight below the surface ($1/\sigma'_t$ is often referred to as the transport mean free path). This is a first-order approximation of the multipole that captures scattering due to a normally incident light ray [Eason et al. 1978]. This displacement of the real source leads to a modified distance to the real source $d_r = \sqrt{r^2 + z_r^2}$ which is based on the assumption that the ray is normally incident on a planar surface [Jensen et al. 2001].

Since our source resembles the actual light ray more than it resembles the first scattering event, we do not displace the real source. This choice introduces a singularity in the model for $r = 0$ (Eq. (9) is valid only for $r \gg 1/\sigma_s$). It also leads to overestimation in the region close to the singularity where the assumption of uniform emergent radiance is invalid. In Appendix B, we explain how to deal with this singularity by comparing our solution to similar ones based on exact transport theory. Although the theoretical reasoning might be complicated, the final solution is quite simple. We correct the distance to the real source using

$$d_r^2 = \begin{cases} r^2 + D\mu_0(D\mu_0 - 2d_e \cos \beta) & , \text{ for } \mu_0 > 0 \text{ (frontlit)} \\ r^2 + 1/(3\sigma_t)^2 & , \text{ otherwise (backlit)} \end{cases} , \quad (24)$$

where $\sigma_t = \sigma_a + \sigma_s$, $\mu_0 = -\vec{n}_o \cdot \vec{\omega}_{12}$, and the other cosine term is

$$\cos \beta = -\sin \theta \frac{r}{\sqrt{r^2 + d_e^2}} = -\sqrt{\frac{r^2 - (\mathbf{x} \cdot \vec{\omega}_{12})^2}{r^2 + d_e^2}} . \quad (25)$$

When introducing a distance correction, we must consider whether the denominator of $\cos \theta$ (appearing in Eqs. (9) and (19)) should be affected by the correction or not. Based on numerical experiments, we have found that correction of the denominator leads to smaller error. We therefore do not explicitly write $\cos \theta$ in the S'_d function (20), as the denominator is then corrected by simply using d_r as the r argument of S'_d for the real source.

3.3.4 Our BSSRDF. Putting it all together, the diffusive part of our BSSRDF is

$$S_d(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o) = S'_d(\mathbf{x}_o - \mathbf{x}_i, \vec{\omega}_{12}, d_r) - S'_d(\mathbf{x}_o - \mathbf{x}_v, \vec{\omega}_v, d_v) , \quad (26)$$

where $\mathbf{x}_v = \mathbf{x}_i + 2Ad_e\vec{n}_i^*$, $d_v = |\mathbf{x}_o - \mathbf{x}_v|$, and $\vec{\omega}_v = \vec{\omega}_{12} - 2(\vec{\omega}_{12} \cdot \vec{n}_i^*)\vec{n}_i^*$ (see Figure 3). To obtain the full BSSRDF, we add the modified reduced intensity term $S_{\delta E}$ and multiply by the Fresnel transmittances T_{12} and T_{21} (Eq. (3)).




Fig. 4. Diffuse reflectance curves for a ray normally incident on planar, nonrefractive semi-infinite media with $\sigma_s = 1 \text{ cm}^{-1}$ but with different absorption properties (σ_a is specified above each plot). The ray is incident at the origin. The graphs plot diffuse reflectance R_d at positions along the x -axis for different models (identified in the legends using the abbreviations specified in Figure 7). Note that each vertical axis uses a log-scale with a different range.

4. IMPLEMENTATION

Just like other BSSRDFs, our model can be used in a number of rendering approaches, the only major difference being that our model takes directions into account. Approaches that assume irradiance as input need relatively 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 and, 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 being 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 set of directions across all irradiance samples in order to spatially cluster them without any change in the original algorithm. Each evaluation of the BSSRDF now goes over the list and takes directions into account. 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 samples across directions. We chose a simpler approach of utilizing the existing implementation of the hierarchical integration method. In the results we present, we used 16 directional samples by importance-sampling the environment map.

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 Section 4. The results in Figures 1, 7, and 8 were rendered using direct Monte Carlo integration. The results in Figure 11 were rendered using hierarchical integration. Optical properties measured by Jensen et

al. [2001] and Narasimhan et al. [2006] were used as input parameters. To validate our model, we compare with full solutions of the radiative transfer equation using unbiased path tracing [Rushmeier 1988]. Single scattering is computed using the same unbiased path-tracing approach, and is added to all the results of the existing models but not to the results of our model.

One common approach to compare BSSRDFs is to plot spatially resolved diffuse reflectance R_d for a ray that is normally incident on a planar semi-infinite medium. In this case, \mathbf{x}_i and $\vec{\omega}_i$ are fixed so that we have

$$R_d(\mathbf{x}_o) = \int_{2\pi} S(\mathbf{x}_i, \vec{\omega}_i; \mathbf{x}_o, \vec{\omega}_o) (\vec{\omega}_o \cdot \vec{n}_o) d\omega_o. \quad (27)$$

Considering a nonrefractive medium with its surface in the xy -plane and a BSSRDF that does not depend on $\vec{\omega}_o$, we have $R_d(x, y) = \pi S_d(x, y)$. Figure 4 shows diffuse reflectance curves for media with different albedos as produced by all the models we tested. We fixed $\sigma_s = 1 \text{ cm}^{-1}$ and tested $\sigma_a = 0.01 \text{ cm}^{-1}$, 0.1 cm^{-1} , 1 cm^{-1} with $\eta_1 = \eta_2 = 1$ and $g = 0$. Our model produces curves that are similar to existing models with slight overestimation. If we take the curves of quantized diffusion and add approximate diffuse single scattering [Habel et al. 2013] instead of path-traced single scattering, we get curves that are very close to ours. This suggests that the overestimation in our model is due to approximation in single scattering. Note that the deviation of our model in the tail of the curve for highly absorbing media ($\sigma_a = 1 \text{ cm}^{-1}$) is in fact very small in its scale. Moreover, as we demonstrate in our other results, the slight deviations in these diffuse reflectance curves are not as significant as the missing directional effects.

To demonstrate the missing directional effects, we also compare diffuse reflectance curves for oblique incidence. Figure 5 uses the same configuration as Figure 4, except that we now change the angle of incidence away from 0° . The x -axis is in the plane of incidence. We tested the medium with $\sigma_a = 0.01 \text{ cm}^{-1}$, and we included the other models to verify that single scattering has minor directional effects. In these plots, only our model captures the characteristic distortion of the curve according to the angle of incidence, especially at 60° . Figure 6 depicts the region around the peak as 2D plots over the surface of the medium. The curves in Figure 5 correspond to horizontal slices through the middle of the plots in Figure 6. In the 2D plots, we compare our model with quantized diffusion plus path-traced single scattering and with unbiased path tracing. While both analytical models do not perfectly match the path-traced reference,




Fig. 5. Diffuse reflectance curves for a ray obliquely incident on a planar, nonrefractive semi-infinite medium. The ray is incident at the origin at different angles with the normal (the angles are specified above each plot), and the x -axis is in the plane of incidence. Curves are plotted for the same models as in Figure 4. We used $\sigma_s = 1 \text{ cm}^{-1}$ and $\sigma_a = 0.01 \text{ cm}^{-1}$.

our model distorts the reflectance according to the angle of incidence as in the reference, especially along the x -axis (parallel to the plane of incidence). The quantized diffusion slices exhibit only very slight variations that are all due to the added single scattering.

Although diffuse reflectance plots for semi-infinite media can be informative, what is more important in practice is the accuracy of rendered images in more complex cases with refractive boundaries that are not flat. Figure 7 compares rendered images of simple shapes with different materials. Figure 9 plots the root-mean-squared error (RMSE) and the structural similarity index (SSIM [Wang et al. 2004]) of the corresponding high dynamic range images (using perceptually uniform encoding for SSIM [Aydın et al. 2008]). Our model produces results that are significantly closer to path traced references when the material is highly forward scattering. For such cases, directionality of incident light is important since light generally maintains its initial direction over longer distances. This makes the directional effects nonnegligible in the resulting radiance distribution. For isotropically scattering materials (first four materials), existing models are comparable to our model. We note, however, that this is true only if accurate path-traced single scattering is added to the existing models (Figure 9, full color). The error nearly doubles if single scattering is not added (Figure 9, faded color).

Figure 1 shows the bunny model rendered with the white grapefruit juice material as an example of more complex geometries. The result of our model is again closer to the path-traced reference. In particular, our model captures bright scattered highlights that are missing in existing models. Figure 8 shows the same scene configuration with the marble material. For this material, existing models tend to work well, yet our model is still more accurate than the other models. The squared differences between BSSRDF results and the path-traced references in Figures 1 and 8 are shown in Figure 10. Figure 11 demonstrates renderings with image-based lighting using hierarchical integration. Our model captures details of the geometry that are highlighted by the directionality of the incident illumination. These details are lost with the existing models.

It is noteworthy that a model which has fairly accurate diffuse reflectance curves does not necessarily produce accurate rendered images. Although our model has normal incidence curves that are slightly less accurate (Figure 4), the rendered images are generally more accurate than if we use the existing models. This discrepancy is due to the facts that the existing models do not consider the direction toward the point of emergence (Figure 6) nor do diffuse reflectance curves capture it (Figure 5), and also that fundamental assumptions such as semi-infinite media are not valid in practice.




Fig. 6. False-color images depicting $R_d(x, y)$ in the region $[-2 \text{ cm}, 2 \text{ cm}] \times [-2 \text{ cm}, 2 \text{ cm}]$. These 2D slices have the same configuration as the plots in Figure 5. The first row is quantized diffusion plus path-traced single scattering (qntzd), the second row is our directional dipole (ours), and the third row is unbiased path tracing (ptrace). The angles of incidence are at the top. We used $\sigma_s = 1 \text{ cm}^{-1}$ and $\sigma_a = 0.01 \text{ cm}^{-1}$. Our model varies significantly with the angle of incidence, while quantized diffusion varies only due to single scattering.

In direct Monte Carlo integration, there is little practical difference in using our model instead of the standard (or better) dipole. For each evaluation of the BSSRDF, we simply replace a single scattering evaluation by a few additional arithmetic operations in the diffusive part. It is significantly faster to evaluate our BSSRDF than to evaluate the quantized diffusion BSSRDF. In order to make the evaluation costs of quantized diffusion of the same order of magnitude as the other models, we needed to implement the radial caching

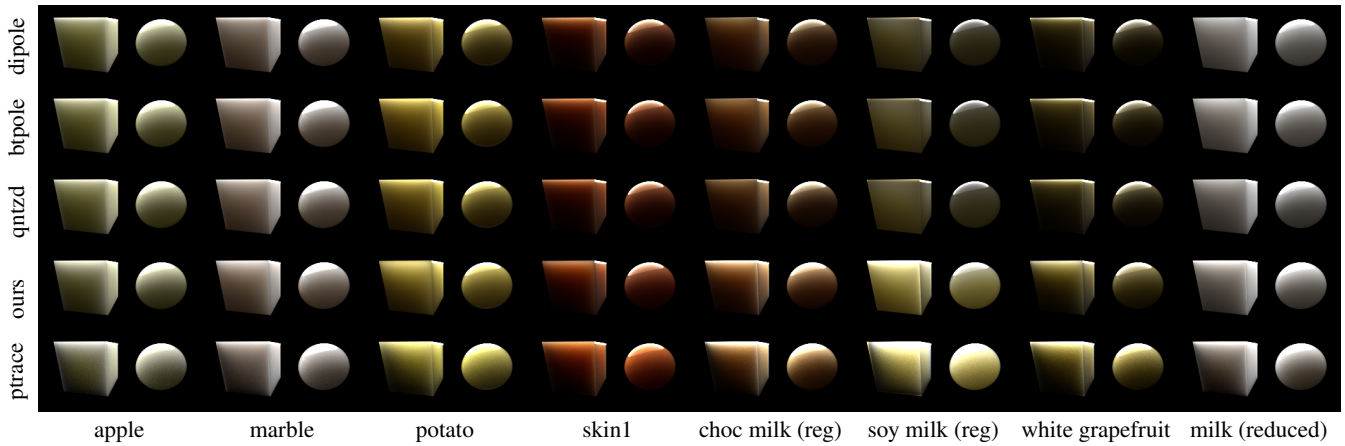


Fig. 7. Simple test cases comparing the standard dipole model (dipole) [Jensen et al. 2001], the better dipole model (btpole) [d’Eon 2012], quantized diffusion (qntzd) [d’Eon and Irving 2011], our model (ours), and the reference solutions rendered by path tracing (ptrace). Path-traced single scattering was added to the existing models but not to ours. Materials used are all from measured values [Jensen et al. 2001; Narasimhan et al. 2006]. The results of our model are closer to the path-traced references or equally goods, while the rendering times are shorter as our model includes single scattering.




Fig. 8. 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 in terms of error. Perceptually, our result seems too smooth, due to the assumption that emerging light is diffuse. The same assumption is used in the other analytical models, but, for those models, the problem is partly mitigated by the path-traced single scattering. The Stanford Bunny model is courtesy of the Stanford University Computer Graphics Laboratory (<http://graphics.stanford.edu/data/3Dscanrep/>).

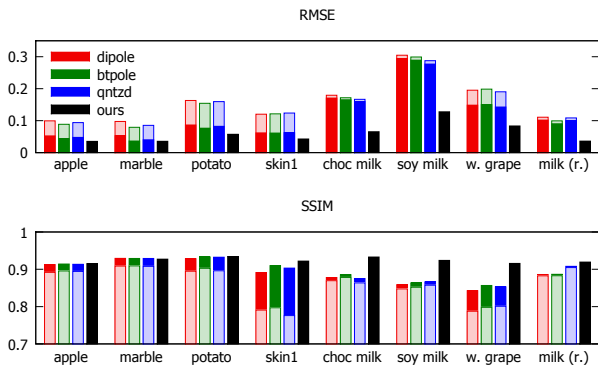


Fig. 9. Objective image quality measurements for the renderings in Figure 7. Every measurement for an existing model (full color) is accompanied by a measurement where path-traced single scattering is excluded (faded color). Our model outperforms the other models, except if materials have high albedo and scatter isotropically, in which case the other models can be equally good when single scattering is added.

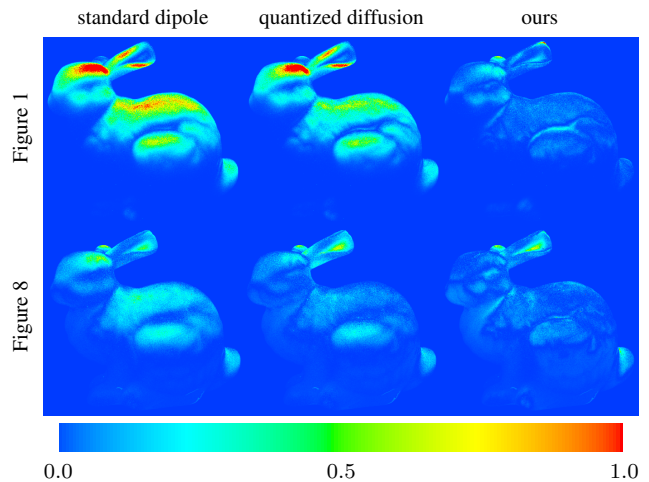


Fig. 10. Squared differences of BSSRDF results versus references.




Fig. 11. Image-based lighting (using the Uffizi Gallery light probe) of the Happy Buddha with the materials also used in the simple test cases (Figure 7). The images were rendered using hierarchical integration of the quantized diffusion model [d’Eon and Irving 2011] plus single scattering added by separate Monte Carlo ray tracing (top row) and our model (bottom row). We excluded Fresnel reflection in order to better highlight the differences between the two models. The Happy Buddha model is courtesy of the Stanford University Computer Graphics Laboratory (<http://graphics.stanford.edu/data/3Dscanrep/>). The Uffizi Gallery light probe is courtesy of Paul Debevec (<http://www.pauldebevec.com/Probes/>).




Fig. 12. Diffuse reflectance curves for a ray incident at a 45° angle without adding path-traced single scattering to the existing models. Both the standard and the better dipole now miss a sharp peak around the origin. While the quantized diffusion model maintains the sharp peak, there is no characteristic shifting of the profile due to the oblique angle of incidence.

technique of d’Eon and Irving [2011]. This technique cannot be used in the same way with our model due to the dependency on directions. Similarly, in our current implementation of hierarchical integration, we must (as opposed to existing models) evaluate the diffusive part of our BSSRDF for every direction of incident illumination that we consider. On the other hand, the single scattering that is added separately to the existing models also depends on directions, and its cost is thus also proportional to the number of directional samples. Usually, the number of directional samples must be very large to obtain noiseless single scattering.

In terms of rendering time, full unbiased path tracing in common scenes with a refractive medium and a small light source (Figures 1 and 8) takes days using one CPU. Excluding multiple scattering in our path-traced results was only a factor of four or less faster. In a scene with a large light source (Figures 7 and 11), path tracing is up to two orders of magnitude faster, but still takes several hundred seconds. It is thus very expensive to add single scattering in the existing models, where it can easily take up 90% of the total rendering time. One option to reduce the computational costs of single scattering is to use some approximation, as suggested in previous work [Hanrahan and Krueger 1993; Jensen et al. 2001; Habel et al. 2013]. Such approximation, however, introduces further error in the rendered images, making our model even more favorable in terms of accuracy. Figure 12 shows the consequences in the diffuse reflectance curves if we exclude single scattering.

In most cases, path traced single scattering makes the total rendering time of existing models larger than if our model is used. However, this is not the case in Figure 11, as we used hierarchical integration and the light source is large (due to the environment map we used). The images in the bottom row of Figure 11, which were rendered with our model using 16 directions and 4 samples per pixel, took 18–20 minutes each on a 3.4 GHz CPU (Intel Core i7-2600) using a single thread. As a representative of the existing models, quantized diffusion with radial caching and with added single scattering produced each of the images in the top row of Figure 11 in 13–14 minutes. Thus, in this unfavorable case of Figure 11, our model is a factor 1.4 more expensive than the existing models.

6. DISCUSSION

6.1 Negative Values

In some cases, our model can return a negative value. Figure 13 illustrates such cases with the index-matched boundary. We have




Fig. 13. Negative values appear in our model at grazing incidence. These 2D slices are similar to those in Figure 6, but with different absorption coefficients and angles of incidence (the values are specified at the bottom). The regions enclosed by a white contour have been clamped to zero. For highly absorbing media, our model can return negative values for even smaller angles. No configuration in Figure 6 has such a clamped region.

found our model tends to return a negative value when the albedo is low and the angle of incidence is large. This is mainly a consequence of the approximations in the ray source solution (9). 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. We also note that negative values rarely occur if a medium has an index of refraction larger than the surrounding medium, as is commonly the case.

6.2 Strictly Planar Model and Modified Tangent Plane

Since our boundary conditions are approximative, as in the existing models, it is possible to consider some other theoretically reasonable variations of the final model in Eq. (26). For example, while our model uses distances and angles computed directly from given locations and directions, we could assume a strictly planar configuration (with \vec{n}_o always perpendicular to \mathbf{x} in Eq. (20)) and compute distances and angles accordingly. This approach is used in existing analytical models [Jensen et al. 2001; d’Eon and Irving 2011; d’Eon 2012]. As explained before, we could also mirror sources according to the normal at the point of incidence (\vec{n}_i) instead of the normal of the modified tangent plane (\vec{n}_i^*).

Figure 14 compares such potential variants of our model. This comparison demonstrates that, if we assume a strictly planar boundary as in the existing models, we would slightly blur illumination features that are present with our current model. Moreover, we have found that the strictly planar variant as well as the variant without the modified tangent plane usually overestimate illumination in backlit regions. Although this result does not conclude that our choices are fundamentally better than other variants, we have found it is rather difficult to improve upon the final model so that accuracy is improved across all our test cases.

6.3 Reciprocity

Unlike the standard dipole model and the quantized diffusion model, both of which only depend on the distance, our model is not reciprocal. In other words, swapping variables for a point of incidence and a point of emergence does not result in the same value in our model. If a reciprocal model is desirable (e.g., for bidirectional rendering algorithms), one can take the average of two evaluations of our model with swapped variables. Figure 15 shows images rendered with this modification. For highly scattering materials (Figure 15, right), we have found our model is close to reciprocal, meaning this modification would not result in significant changes in rendered im-




Fig. 14. Comparison of our current model, our model if we assume a strictly planar boundary, and our model if we use the tangent plane at the point of incidence instead of the modified tangent plane. The images are closeups of the bunny in Figure 8 without Fresnel reflections to highlight the differences. The strictly planar variant is close to the current model, but misses subtle details and slightly overestimates. The use of the modified tangent plane is important for our model to work well.




Fig. 15. Taking the average of two evaluations of our model with the variables of incoming and outgoing light swapped around would make our model reciprocal. The difference from the nonreciprocal image in Figure 11 depends on the material.

ages. For thinner media (left), however, this modification produces some differences in rendered images. We have not fully investigated whether these differences are generally desirable or not in terms of accuracy. All other results in this article do not use this modification.

6.4 Limitations

Since we still make a number of assumptions to derive our model, it does not perfectly match the results of path tracing (this is evident in Figure 6, for example). Our model also does not reproduce emergent radiance distributions that are not uniform. In addition, our model does not distinguish between materials with different scattering coefficients (σ_s) and asymmetry parameters (g) if they result in the same reduced scattering coefficient ($\sigma'_s = \sigma_s(1 - g)$), except in the correction of the distance to the real source (Section 3.3.3).

We emphasize again that, similar to other models, ours is derived under the assumption of a planar semi-infinite medium. The modified tangent plane does not remove this limitation. This assumption is known to have issues in a number of cases, such as sharp boundaries and thin features. Extensions of our model to quadpole and multipole constructions to alleviate these issues are interesting future work.

7. CONCLUSION

We introduced a novel BSSRDF that comes closer to path-traced references than existing models. Our model is built upon a new analytical solution to the diffusion equation that enables us to take the direction of incoming light into account. We explained how to derive a practical BSSRDF by extending this solution into a dipole of real and virtual ray sources. Our work significantly improves

the accuracy of an analytical model based on diffusion theory. Improvement over existing models is especially significant for highly forward scattering materials, where the directionality of the incident illumination is more important. We believe that our model is a valuable addition to the existing set of analytical BSSRDFs. Since our model is fully analytic, application to real-time rendering is very promising as compared with models based on numerical integration.

APPENDIX

A. FRESNEL INTEGRALS

In Section 3.2, we need hemispherical integrals of the Fresnel transmittance T . More specifically, we need to evaluate the cosine-weighted hemispherical integral of T_{21} multiplied by L_d in Eq. (14). To do this, we use the functions

$$C_\phi(\eta) = \frac{1}{4\pi} \int_{2\pi} T_{21}(\eta, \theta_o) \cos \theta_o \, d\omega_o \quad (28)$$

$$C_E(\eta) = \frac{3}{4\pi} \int_{2\pi} T_{21}(\eta, \theta_o) \cos^2 \theta_o \, d\omega_o, \quad (29)$$

where $\cos \theta_o = \vec{n}_o \cdot \vec{\omega}_o$. In the following, we leave out the argument of the C -functions to shorten notation. The C -functions are easily written in terms of Fresnel reflectance R instead of Fresnel transmittance, since $T = 1 - R$. This means that

$$C_\phi = \frac{1}{4\pi} \left(\pi - \int_{2\pi} R_{21}(\eta, \theta_o) \cos \theta_o \, d\omega_o \right) = \frac{1}{4} (1 - 2C_1) \quad (30)$$

$$C_E = \frac{3}{4\pi} \left(\frac{2\pi}{3} - \int_{2\pi} R_{21}(\eta, \theta_o) \cos^2 \theta_o \, d\omega_o \right) = \frac{1}{2} (1 - 3C_2). \quad (31)$$

D'Eon and Irving [2011] also use the functions C_ϕ and C_E , and they provide the following convenient polynomial approximations to ease their evaluation:

$$2C_1 \approx \begin{cases} 0.919317 - 3.4793\eta + 6.75335\eta^2 - 7.80989\eta^3 \\ \quad + 4.98554\eta^4 - 1.36881\eta^5, & \eta < 1 \\ -9.23372 + 22.2272\eta - 20.9292\eta^2 + 10.2291\eta^3 \\ \quad - 2.54396\eta^4 + 0.254913\eta^5, & \eta \geq 1 \end{cases}$$

$$3C_2 \approx \begin{cases} 0.828421 - 2.62051\eta + 3.36231\eta^2 - 1.95284\eta^3 \\ \quad + 0.236494\eta^4 + 0.145787\eta^5, & \eta < 1 \\ -1641.1 + \frac{135.926}{\eta^3} - \frac{656.175}{\eta^2} + \frac{1376.53}{\eta} + 1213.67\eta \\ \quad - 568.556\eta^2 + 164.798\eta^3 - 27.0181\eta^4 + 1.91826\eta^5, & \eta \geq 1. \end{cases}$$

B. DISTANCE CORRECTION

Our distance correction is inspired by more general versions of the standard dipole and the relation between these diffusion dipoles and similar solutions from more exact transport theory. The main idea is to see how our model without the distance correction differs in terms of the distance between the points of incidence and emergence. We limit our investigation to the planar semi-infinite medium, which is also the case that we considered when configuring the dipole.

Suppose we orient our coordinate system such that the xy -plane is the surface and the ray is incident at the origin with the z -axis along $-\vec{n}_o$. Assuming that the medium exhibits isotropic scattering and that the solution is independent of outgoing direction, the emergent radiance $L_r(\mathbf{x}_o, \vec{\omega}_o)$ becomes an integral over z from 0 to ∞ of a

Hankel transform of a Green function which is Fourier transformed in the x and y coordinates [Williams 1982]. Analytical evaluation of a diffuse reflectance integral like this is difficult except in the special cases of a point source [Elliott 1955] and a ray source in the normal direction [Williams 2009]. In these two ‘‘monopole’’ solutions that incorporate a planar boundary, an approximate analytical evaluation of the diffuse reflectance integral is possible if the radial distance is corrected to

$$R^2 = r^2 + (z' + d_e)^2, \quad (32)$$

where z' is the depth of the point source or the internal ray source. For $z' = 0$, Elliott [1955] observed that the d_e^2 term in Eq. (32) essentially deals with the boundary in a way somewhat similar to what the standard dipole does. The main difference is that radial distance is corrected in the transport solution, whereas in the standard dipole the diffuse reflectance integral is approximated by displacing the real source ($z' = z_r$) and mirroring it in an extrapolated boundary to place the virtual source (instead of d_e^2). Our approach is motivated by this observation; we retain the extrapolated boundary by the virtual source, but we also introduce a radial distance correction for the real source.

It has been observed in work based on the diffusion dipole that the distance correction should change if the point of emergence is not in the tangent plane [Fretterd and Longini 1973], or if the ray is incident at an oblique angle [Wang and Jacques 1995]. The radial distance correction from transport theory (Eq. (32)) also needs modification in these cases, since z' and d_e are no longer taken in the same direction. Then, considering triangle geometry,

$$R^2 = r^2 + z'^2 + d_e^2 - 2z'd_e \cos \beta, \quad (33)$$

where we approximate $\cos \beta$ as in Eq. (25).

Considering Elliott’s comparison of the transport solution and the diffusion dipole, we do not include the d_e^2 term in our radial distance correction as this part is included with the virtual source. At oblique incidence, the diffuse reflectance integral can be reformulated as an integral along the refracted ray. This integral can be approximated to first order by choosing an offset D^* along the refracted ray. Then

$$z' = D^* |\mu_0| \quad \text{with} \quad \mu_0 = -\vec{\omega}_{12} \cdot \vec{n}_o, \quad (34)$$

where $D^* = D$ if the surface is frontlit ($\mu_0 > 0$) and $D^* = 1/(3\sigma_t)$ if the surface is backlit ($\mu_0 \leq 0$). The use of D^* instead of $3D$ was determined empirically. Our corrected radial distance for the real source is then Eq. (24), where we in the backlit case (due to some minor issues along sharp edges) decided to set the correction term as if light were normally incident with \mathbf{x}_o located along the normal direction (such that $\mu_0 = -1$ and $\cos \beta = \sin \theta = 0$).

ACKNOWLEDGMENTS

We would like to thank Jesper Mosegaard for useful discussions. 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.
- AYDIN, T. O., MANTIUK, R., AND SEIDEL, H.-P. 2008. Extending quality metrics to full luminance range images. In *Human Vision and Electronic Imaging XIII*. Proceedings of SPIE, vol. 6806. 68060B.
- BEVILACQUA, F. AND DEPEURSINGE, C. 1999. Monte Carlo study of diffuse reflectance at source-detector separations close to one transport mean free path. *Journal of the Optical Society of America A* 16, 12 (December), 2935–2945.

- BLINN, J. F. 1982. Light reflection functions for simulation of clouds and dusty surfaces. *Computer Graphics (Proceedings of ACM SIGGRAPH 82)* 16, 3 (July), 21–29.
- DAVISON, B. 1958. *Neutron Transport Theory*. Oxford University Press. Authored with the collaboration of J. B. Sykes.
- D'EON, E. 2012. A better dipole. Publicly available technical report. <http://www.eugenedeon.com/?project=a-better-dipole>.
- 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.
- 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.
- DORSEY, J., EDELMAN, A., JENSEN, H. W., LEGAKIS, J., AND PEDERSEN, H. K. 1999. Modeling and rendering of weathered stone. In *Proceedings of ACM SIGGRAPH 1999*. 225–234.
- EASON, G., VEITCH, A. R., NISBET, R. M., AND TURNBULL, F. W. 1978. The theory of the back-scattering of light by blood. *Journal of Physics D: Applied Physics* 1978, 10 (July), 1463–1479.
- ELLIOTT, J. P. 1955. Milne's problem with a point-source. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 228, 1174 (March), 424–433.
- 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 optical 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.
- FRETTERD, R. J. AND LONGINI, R. L. 1973. Diffusion dipole source. *Journal of the Optical Society of America* 63, 3 (March), 336–337.
- GLASTONE, 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.
- HABEL, R., CHRISTENSEN, P. H., AND JAROSZ, W. 2013. Photon beam diffusion: A hybrid Monte Carlo method for subsurface scattering. *Computer Graphics Forum (Proceedings of EGSR 2013)* 32, 4 (June), 27–37.
- HANRAHAN, P. AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Computer Graphics (Proceedings of ACM SIGGRAPH 93)*, 165–174.
- 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. AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of ACM SIGGRAPH 98*. 311–320.
- 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.
- JOSEPH, J. H., WISCOMBE, W. J., AND WEINMAN, J. A. 1976. The delta-Eddington approximation for radiative flux transfer. *Journal of Atmospheric Sciences* 33, 2452–2459.
- KIENLE, A. 2005. Light diffusion through a turbid parallelepiped. *Journal of the Optical Society of America A* 22, 9 (September), 1883–1888.
- 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.
- MENON, S., SU, Q., AND GROBE, R. 2005a. Determination of g and μ using multiply scattered light in turbid media. *Physical Review Letters* 94, 153904.
- MENON, S., SU, Q., AND GROBE, R. 2005b. Generalized diffusion solution for light scattering from anisotropic sources. *Optics Letters* 30, 12 (June), 1542–1544.
- NARASIMHAN, S. G., GUPTA, M., DONNER, C., RAMAMOORTHI, 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 LIMPERS, 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. Ph.D. thesis, Cornell University.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques '95*, P. Hanrahan and W. Purgathofer, Eds. Springer, 41–50. Proceedings of the Sixth Eurographics Workshop on Rendering.
- SUN, B., RAMAMOORTHI, R., NARASIMHAN, S. G., AND NAYAR, S. K. 2005. A practical analytic single scattering model for real time rendering. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2005)* 24, 3, 1040–1049.
- WALTER, B., ZHAO, S., HOLZSCHUCH, N., AND BALA, K. 2009. Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2009)* 28, 3 (August), 92:1–92:8.
- WANG, L. AND JACQUES, S. L. 1995. Use of a laser beam with an oblique angle of incidence to measure the reduced scattering coefficient of a turbid medium. *Applied Optics* 34, 13 (May), 2362–2366.
- 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.
- 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.
- WILLIAMS, M. M. R. 1982. The three-dimensional transport equation with applications to energy deposition and reflection. *Journal of Physics A: Mathematical and General* 15, 3 (March), 965–983.
- WILLIAMS, M. M. R. 2009. Three-dimensional transport theory: An analytical solution of an internal beam searchlight problem-I. *Annals of Nuclear Energy* 36, 6 (June), 767–783.
- 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.

Received December 2013; revised May 2014; accepted May 2014