

A radiative transfer framework for rendering materials with anisotropic structure

Wenzel Jakob

Adam Arbree

Jonathan T. Moon
Cornell University

Kavita Bala

Steve Marschner




Figure 1: Renderings of a scarf represented as a high-resolution volume. Accounting for the anisotropic structure of the medium leads to a significantly changed appearance, including realistic highlights and color variations.

Abstract

The radiative transfer framework that underlies all current rendering of volumes is limited to scattering media whose properties are invariant to rotation. Many systems allow for “anisotropic scattering,” in the sense that scattered intensity depends on the scattering angle, but the standard equation assumes that the structure of the medium is isotropic. This limitation impedes physics-based rendering of volume models of cloth, hair, skin, and other important volumetric or translucent materials that do have anisotropic structure. This paper presents an end-to-end formulation of physics-based volume rendering of anisotropic scattering structures, allowing these materials to become full participants in global illumination simulations.

We begin with a generalized radiative transfer equation, derived from scattering by oriented non-spherical particles. Within this framework, we propose a new volume scattering model analogous to the well-known family of microfacet surface reflection models; we derive an anisotropic diffusion approximation, including the weak form required for finite element solution and a way to compute the diffusion matrix from the parameters of the scattering model; and we also derive a new anisotropic dipole BSSRDF for anisotropic translucent materials. We demonstrate results from Monte Carlo, finite element, and dipole simulations. All these contributions are readily implemented in existing rendering systems for volumes and translucent materials, and they all reduce to the standard practice in the isotropic case.

ACM Reference Format

Jakob, W., Arbree, A., Moon, J., Bala, K., Marschner, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.* 29, 4, Article 53 (July 2010), 13 pages. DOI = 10.1145/1778765.1778790 <http://doi.acm.org/10.1145/1778765.1778790>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 0730-0301/2010/07-ART53 \$10.00 DOI 10.1145/1778765.1778790 <http://doi.acm.org/10.1145/1778765.1778790>

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Three-Dimensional Graphics and Realism

Keywords: Subsurface scattering, anisotropy, light transport, diffusion theory, finite element method, dipole model, BSSRDF.

1 Introduction

Volume models are valuable in rendering materials with complex surface or volumetric structure [Perlin and Hoffert 1989; Kajiya and Kay 1989; Xu et al. 2001; Neyret 1998], and many very successful projects have made detailed visualizations and realistic renderings using them. However, there is a fundamental gap in the technology for rendering volumes, which impedes progress in improving the realism of volume renderings. On one hand, the tradition of volume visualization, or “volume rendering,” achieves high detail and realistic shading effects using heuristic shading models that can describe directional scattering from surfaces and fibers, but these models lack a sound physical basis. On the other hand, physically based rendering systems support volumes as “participating media” using physically sound formulations originating from atmospheric optics; unfortunately, these models are limited to isotropic media¹ and are inherently unable to describe the directional scattering needed to achieve realistic appearance in nontrivial solid materials.

This paper aims to bridge this gap by upgrading the underlying

¹We use the word “isotropic” to refer to a medium with properties that are invariant to rotating the medium, in contrast to the alternative meaning of scattering that is independent of scattered direction.




Figure 2: The distinction between isotropic and anisotropic media.

framework of participating media to handle anisotropic effects, such as specular reflections from surfaces or fibers. This allows the machinery of physics-based rendering to be brought to bear on a wider range of volume models, thus providing the foundation for a new level of realism in volume models. Currently, participating media rendering algorithms model scattering media using a radiative transfer model limited to isotropic media—media whose properties are invariant to rotation. For instance, a cloud of steam filled with spherical water droplets fits within the standard model; or a cirrus cloud filled with randomly oriented ice crystals: either medium looks the same to light traveling in different directions. But a volume filled with aligned ice crystals; or hair, wood, or cloth fibers; or fibrous tissue like muscle or collagen; does not fit the usual model because its properties depend on orientation. This is the type of medium we call *anisotropic* in this paper (Figure 2). (This is a deviation from the common usage in graphics, where “anisotropic” refers to scattering that depends on the scattering angle, such as the strong forward scattering exhibited by spherical water drops.)

Anisotropic media have fundamentally different behavior from isotropic media, but they can be handled without burdensome changes to existing renderers. This paper outlines the changes required to the theory and implementation of renderers for participating media and translucent materials, to make them handle anisotropic media. Specifically, we provide:

- (i) an anisotropic version of the radiative transport equation (RTE) to handle anisotropic media, derived from scattering by oriented non-spherical particles;
- (ii) a new model for oriented phase functions that generalizes existing phase functions while also representing surface-like and fiber-like scattering;
- (iii) an anisotropic generalization of the widely used diffusion approximation for rendering translucent materials; and
- (iv) methods for computing light transport solutions using Monte Carlo to directly solve the RTE, using finite elements to solve the diffusion approximation in a translucent material, or using a new anisotropic version of the widely used dipole diffusion solution.

These contributions establish the framework required for further development of volume models as first-class objects in modern physics-based rendering systems.

2 Prior work

Radiative transfer is used in many areas including astrophysics, hydrology, neutron transport, remote sensing, acoustics, seismology, and medical physics, and the standard form of the RTE comes from references in these fields [van de Hulst 1957; Chandrasekhar 1959; Preisendorfer 1976; Ishimaru 1978; Mishchenko et al. 2006]. The theory is usually stated for the prevailing case of spherical or randomly oriented particles, disallowing anisotropic media (but allowing angle-dependent scattering), and this form has generally been used in graphics.

Research in volume rendering has introduced a variety of methods for mapping volume data into an image, including ray tracing [Levoy 1988], resampling [Drebin et al. 1988; Lacroute and Levoy 1994], and rasterization [Westover 1990] approaches, with many recent systems relying on GPU ray casting [Crassin et al. 2009]. While the standard ray integration equation comes from the RTE, these systems do not solve the full equation and instead use heuristic shading methods derived from surface reflection models, often with the density gradient playing the role of the surface normal.

The development of physics-based surface renderers has seen most rendering algorithms extended to handle participating media, from early papers on ray tracing [Kajiya and Herzen 1984; Blinn 1982] and radiosity [Rushmeier and Torrance 1987] to subsequent global illumination systems [Lafortune and Willems 1996; Jensen and Christensen 1998; Pauly et al. 2000; Cerezo et al. 2005]. Most often, participating media are used to add general fog or haze to scenes, but they are also used to make realistic images of detailed atmospheric phenomena like smoke [Fedkiw et al. 2001], and clouds [Max 1994]. Many papers have investigated efficient real-time rendering methods based on analytic solutions in special cases [Premoze et al. 2004; Sun et al. 2005]. Because of the goal of physical accuracy and the need to integrate with global illumination systems, all these methods pay attention to correct implementation of the RTE, inheriting the important restriction to isotropic media.

Between these two camps lie techniques that use volume rendering to make realistic images of materials, such as fur [Kajiya and Kay 1989], cloth [Xu et al. 2001], and other materials [Perlin and Hoffert 1989], that are difficult to represent with surfaces. These models must either stick with isotropic scattering or accept being incompatible with the physics-based-rendering toolbox—both undesirable choices when realism is the goal. Some papers, notably the volume modeling system of Neyret [1998], have recognized and discussed many of the relevant effects, but have not made the connection to physically based models of volume scattering that is required for these models to be used in full physics-based simulations.

Many models for scattering from surfaces or in volumes have been proposed, with microfacet models for surfaces [Cook and Torrance 1982; Ashikhmin et al. 2000; Walter et al. 2007] and the Henyey-Greenstein model for volumes [Cerezo et al. 2005] being the most widely used. The model proposed here has similarities to microfacet models and to Neyret’s volumetric textures [1998], but differs in that it provides a physically based model for the volume setting.

The diffusion approximation to radiative transfer is widely used to approximate transport of light or other energy in highly scattering media. It is derived from the RTE either using a first-order approximation to the directional radiance [Ishimaru 1978], as we do in this paper, or by taking a limit of diminishing mean free path [Ryzhik et al. 1996]. In graphics, diffusion is used to describe light transport through translucent materials [Jensen et al. 2001]. Previous rendering work has always used isotropic diffusion, but other fields have developed methods for anisotropic diffusion. Because many human tissues exhibit significant anisotropy, recent work in medical physics has examined the effects of anisotropy [Kienle et al. 2004]. In particular, optical tomography [Arridge 1999; Gibson et al. 2005] uses the diffusion approximation inversely to recover volume structure from external observations, and recent works [Heiskala et al. 2005; Heino et al. 2003] have found anisotropic diffusion approximations similar to ours by replacing individual components of the isotropic RTE with more expressive counterparts. However, these approaches do not account for the subtle interrelations between components of the RTE and complications regarding reciprocity and normalization in anisotropic media. Both models also lack expressiveness when rendering directly using the RTE, being limited to isotropic scattering or unable to handle distributions of particles, respectively. These issues motivated the derivation of the anisotropic RTE presented in this paper.

Dipole sources can provide simple approximate solutions when diffusion occurs in a homogeneous region with a planar boundary. This is the prevailing way to handle translucency in rendering [Jensen et al. 2001], and similar formulations can be found in other fields. In geoscience, dipole solutions have been used for transport of electric charge in the earth, and in that context an anisotropic dipole was proposed [Li and Uren 1998], which inspired our anisotropic BSSRDF

model. In practice, many differences between the two models exist due to the different underlying equations they solve.

In biophysics, recent work has led to both time-resolved [Dudko and Weiss 2005] and steady-state [Johnson and Legendijk 2009] anisotropic multipole models for slab geometries. Both models use absorptive boundary conditions that are inappropriate for graphics applications, since they force the average illumination to zero on the boundary. Another issue is that the supplied solution for the radiance on the slab boundaries only applies to the transmission mode, whereas reflection is generally needed for rendering. Our model provides an expression for this direction and, as will be seen later, can also easily be extended to a multipole configuration.

Johnson and Legendijk's [2009] work presented an approach to relate micro- and macroscale anisotropy in forming their anisotropic multipole model. Because they consider only the shape of the modeled particles, but not their scattering properties, the final diffusion matrix derives its anisotropy solely from the particles' directionally varying projected area. In comparison, we present a formulation that accounts for anisotropy in the phase function, which leads to an additional integral term in the resulting diffusion matrix.

The contribution of this paper is to upgrade the framework of participating media rendering so that it works for anisotropic media, which makes it applicable to volume models of the sort that could previously only be rendered using less-principled volume rendering methods. This involves generalizing the RTE, the diffusion approximation, and the dipole BSSRDF model—each by employing methods similar to methods found in the physics, medical imaging, and geosciences literature, but each requiring changes to become suitable for use in rendering systems. Rendering anisotropic media also requires expressive and detailed models for anisotropic scattering, which have not been developed in other fields, leading us to propose the new anisotropic scattering model in section 4.

3 Anisotropic Radiative Transfer

The isotropic radiative transfer equation (RTE), in the form customarily used in graphics, is

$$(\omega \cdot \nabla)L(\omega) = -\sigma_t L(\omega) + \sigma_s \int_{S^2} f_p(\omega' \cdot \omega) L(\omega') d\omega' + Q(\omega). \quad (1)$$

(In this equation all quantities depend on position, so this dependence is left implicit for compactness.) This equation relates the directional derivative of the radiance distribution L to its value via the attenuation coefficient σ_t and to an integral of L at the same point via the scattering coefficient σ_s and the phase function f_p . Energy comes into the equation via the volume source term Q . This equation is for an isotropic² medium because the interaction coefficients σ_s and σ_t are independent of direction, and because the phase function depends only on the angle between ω and ω' .

In order for this equation to describe an anisotropic medium, f_p must be a function of two directions independently, rather than of a single angle: $f_p(\omega \rightarrow \omega')$. Furthermore σ_s and σ_t also depend on ω , leading to the anisotropic RTE:

$$(\omega \cdot \nabla)L(\omega) + \sigma_t(\omega)L(\omega) = \sigma_s(\omega) \int_{S^2} f_p(\omega' \rightarrow \omega) L(\omega') d\omega' + Q(\omega). \quad (2)$$

Although this generalization is superficially straightforward, subtle complications with normalization and reciprocity must be resolved

²Again, we mean “invariant to rotation,” not “independent of direction.”

before we can define a valid phase function model or derive a diffusion approximation. In the isotropic RTE, f_p is a probability distribution (when either argument is held fixed) and reciprocity requires that it be invariant to exchanging its arguments. In the anisotropic case we must choose between keeping f_p normalized as a probability distribution and keeping it symmetrical; both cannot be achieved at once.

For example, consider a volume containing a distribution of aligned needle-shaped particles with high albedo. Such a volume has much lower σ_t (and σ_s) for light traveling parallel to the needles than perpendicular to them. If a volume of these particles is illuminated from the end and viewed from the side, then illuminated from the side and viewed from the end, reciprocity demands that the total scattering remain the same. But since the amount of light interacting with the particles is very different, the phase function values must likewise undergo large changes. Reciprocity still holds in the physical system—its manifestation in the equation is just different. With a somewhat different definition, the phase function can be made symmetrical, but at the expense of no longer being a probability distribution. Again, energy conservation still holds in the physical system, but its manifestation in the equation is different. We have opted for a normalized f_p because it is convenient in the context of Monte Carlo rendering, and the equation is simpler and more familiar.

The isotropic RTE can be derived straightforwardly by analyzing the behavior of a volume filled with spherical scattering particles (spherical because their properties are invariant to rotation). Here we follow this derivation in the style of Ishimaru [1978] but with non-spherical particles, leading to the anisotropic RTE. In both cases it's important to recognize that the resulting equation can also be used successfully to model media that are not made up of particles—biological tissue, cloth, hair, foam, etc.

3.1 Isolated non-spherical particle

Consider an isolated particle illuminated by incident radiance $L(\omega)$. We can characterize the particle using three functions:

- (i) $\sigma(\omega)$ is the area of the particle's projection onto ω^\perp .
- (ii) $0 \leq \alpha(\omega) \leq 1$ is the albedo of the particle when illuminated from direction ω .
- (iii) $p(\omega \rightarrow \omega')$ is the phase function exhibited by the particle when illuminated from direction ω and forms a probability density in the outgoing direction ω' .

When the particle receives collimated light in direction ω , $\sigma(\omega)$ is proportional to the probability of it hitting the particle; $\alpha(\omega)$ is the probability that light is scattered rather than absorbed, conditioned on having hit the particle; and $p(\omega \rightarrow \omega')$ is the probability density for scattering to direction ω' , conditioned on having scattered. Note the distinction between the functions f_p and p : whereas f_p describes the overall behavior of the statistical medium, p refers to the properties of one of its constituent particles.

When the particle is illuminated by a radiance distribution $L(\omega)$, the scattered intensity is found by integrating the scattered light over incoming directions:

$$I(\omega) = \int_{S^2} p(\omega' \rightarrow \omega) \alpha(\omega') \sigma(\omega') L(\omega') d\omega'. \quad (3)$$

Unlike in the case of a spherical particle, for which σ and α are constant, $p(\omega \rightarrow \omega') \neq p(\omega' \rightarrow \omega)$ for a non-spherical particle, and p is generally *not* a probability density as a function of its first argument. The appropriate reciprocity principle is that the scattered intensity to a single direction under collimated illumination remains

invariant under an interchange of the source and receiver. This causes p to obey the more involved reciprocity relation

$$p(\omega \rightarrow \omega')\alpha(\omega)\sigma(\omega) = p(\omega' \rightarrow \omega)\alpha(\omega')\sigma(\omega'). \quad (4)$$

It is the whole chain of events, not just the choice of scattered direction, that has to occur with a reciprocal probability.

3.2 Distribution of particles in a volume

The properties of a volume containing many scattering particles depend on the characteristics of the particles—both their individual properties and the distribution of particle orientations that is present. Up to now there has been no restriction on the shape of particles, but at this point, because it suffices for the model proposed in the next section, we assume that particles are rotationally symmetric about some axis, so that their orientation can be entirely described by the direction of the axis. For the same reason we further assume that the particles are identical.³

Under these two assumptions we can characterize the volume in the neighborhood of a particular point by two quantities:

- (i) ρ is the density of particles per unit volume.
- (ii) $D(m)$, a probability density on the sphere, gives the probability for a particle to be oriented in direction m .

We also define $\sigma(m, \omega)$, $\alpha(m, \omega)$, $p(m, \omega \rightarrow \omega')$, and $I(m, \omega)$ to be the properties of a particle with orientation m .

As in the isotropic case, the effect on light traveling along a ray can be seen by thinking of a cylindrical beam of cross-sectional area A (Figure 3) along which radiance L propagates in the direction ω . An infinitesimal segment ds of the beam contains $n = A ds \rho$ particles, each of projected area $\sigma(m, \omega)$.

Since they are small and well separated, they do not shadow one another, and together they intercept an intensity equal to the radiance times their combined area:

$$\frac{dI_{\text{atten}}(\omega)}{ds} = A \left[\rho \int_{S^2} \sigma(m, \omega) D(m) dm \right] L(\omega).$$

As in derivations of the isotropic RTE, the previous assumption allows particles to be treated independently, and is presumed to hold in the limit of increasingly smaller cylindrical regions. The integral over m accounts for the presence of particles with a distribution of orientations. At the same time, the particles scatter intensity into the direction ω proportional to their number times the intensity (3) averaged over orientations:

$$\frac{dI_{\text{in scat}}(\omega)}{ds} = A \left[\rho \int_{S^2} I(m, \omega) D(m) dm \right].$$

In the limit, these changes in intensity for our finite volume define the directional derivative of radiance (details in supplementary material):

$$(\omega \cdot \nabla)L(\omega) = -dI_{\text{atten}}(\omega)/(A ds) + dI_{\text{in scat}}(\omega)/(A ds)$$

which can be rearranged as

$$(\omega \cdot \nabla)L(\omega) + \sigma_t(\omega)L(\omega) = \int_{S^2} f_a(\omega' \leftrightarrow \omega)L(\omega') d\omega' \quad (5)$$

³Retaining the possibility of more asymmetric particles would require generalizing the ensuing derivation to a distribution over unit quaternions rather than over unit 3-vectors. A mixture of particle types can be accommodated easily by summing or integrating over the particles' properties.




Figure 3: The effect of attenuation and scattering on radiance traveling through a volume of particles with a distribution of orientations.

where $\sigma_t(\omega)$ is the anisotropic volume attenuation coefficient:

$$\sigma_t(\omega) = \rho \int_{S^2} \sigma(m, \omega) D(m) dm \quad (6)$$

and f_a is an inscattering weighting function:

$$f_a(\omega' \leftrightarrow \omega) = \rho \int_{S^2} p(m, \omega' \rightarrow \omega)\alpha(m, \omega')\sigma(m, \omega')D(m) dm.$$

To express f_a in a more convenient form, we define the anisotropic volume scattering coefficient $\sigma_s(\omega)$ as the inscattering due to uniform unit radiance, or $\int f_a(\omega' \leftrightarrow \omega) d\omega'$, which works out to:

$$\sigma_s(\omega) = \rho \int_{S^2} \alpha(m, \omega)\sigma(m, \omega)D(m) dm. \quad (7)$$

Here we have made use of the reciprocity relation (4) and the normalization of p over the outgoing direction. We then divide f_a by σ_s , thereby normalizing it over the incoming direction, to obtain the phase function f_p :

$$f_p(\omega' \rightarrow \omega) = \frac{f_a(\omega' \leftrightarrow \omega)}{\sigma_s(\omega)} = \frac{\rho}{\sigma_s(\omega)} \int_{S^2} p(m, \omega' \rightarrow \omega)\alpha(m, \omega')\sigma(m, \omega')D(m) dm. \quad (8)$$

which inherits from f_a the reciprocity property:

$$\sigma_s(\omega)f_p(\omega' \rightarrow \omega) = \sigma_s(\omega')f_p(\omega \rightarrow \omega'). \quad (9)$$

Substituting $\sigma_s(\omega)f_p(\omega \rightarrow \omega')$ for $f_a(\omega \leftrightarrow \omega')$ in (5), and adding the source term Q , results in the desired form (2).

Together the expressions (6), (7), and (8) define the anisotropic RTE based on the properties of the non-spherical particles that fill the medium. In the next section we examine a particular type of simple particle that can be used to define a phase function model encompassing a range of scattering behaviors.

4 Micro-flake model

We consider a model based on two-sided specularly reflecting flakes oriented according to a known directional distribution. We call this a *micro-flake model*. There are obvious parallels to the popular family of microfacet models for surface reflection [Cook and Torrance 1982]: as with microfacet models, the flakes are not individually resolved but instead aggregated statistically to capture the overall optical behavior of the medium. Another similarity is the use of a half-vector formulation, which follows naturally from the description of a specular surface. The freedom of choosing the micro-flake distribution will later allow the description of various types of anisotropic effects in the context of volume scattering.

We assume that the flakes are planar, and that the albedo of a single flake depends only on the inclination of incident illumination. Using the framework of section 3 to find the phase function for this type of particle entails defining the appropriate functions σ , α and p .

First, due to their planarity, the projected area of a flake with normal m onto ω^\perp is equal to

$$\sigma(m, \omega) = a |\omega \cdot m|$$

where a is the surface area of one side. The above restriction on the albedo of a single particle translates into $\alpha(m, \omega) = \alpha(|\omega \cdot m|)$. That is, α is an arbitrary function of the angle between ω and m that does not distinguish between sides of a flake. The flakes are ideal two-sided reflectors, causing the phase function of a single flake to be a sum of two delta functions associated with the measure $d\omega$:

$$p(m, \omega' \rightarrow \omega) = \delta_\omega(m, h(\omega, -\omega')) + \delta_\omega(-m, h(\omega, -\omega')) \quad (10)$$

Here, $h(\omega_i, \omega_o) := (\omega_i + \omega_o) / \|\omega_i + \omega_o\|$ denotes the half-direction vector. Note the sign on ω' —this direction needs to be flipped due to the common convention that the incident direction of a phase function points inward. Normalization follows since only one of the delta functions becomes nonzero for fixed incoming ω' and m .

To express the delta functions in (10) with respect to particle orientations m and their associated solid angle measure, a correction factor must be applied [Walter et al. 2007]:

$$p(m, \omega' \rightarrow \omega) = \frac{\delta_m(m, h(\omega, -\omega')) + \delta_m(m, -h(\omega, -\omega'))}{4|\omega \cdot m|}$$

and this function satisfies the reciprocity criterion (4), as light can only be scattered when $m \cdot \omega = m \cdot \omega'$. Applying the tools from Section 3 then leads to the following expressions for σ_t , σ_s and f_p (details in supplementary material):

$$\sigma_t(\omega) = a\rho \int_{S^2} |\omega \cdot m| D(m) dm \quad (11)$$

$$\sigma_s(\omega) = a\rho \int_{S^2} \alpha(|\omega \cdot m|) |\omega \cdot m| D(m) dm \quad (12)$$

$$f_p(\omega' \rightarrow \omega) = \frac{a\rho}{4\sigma_s(\omega)} \alpha(|\omega' \cdot h(\omega, -\omega')|) (D(h(\omega, -\omega')) + D(-h(\omega, -\omega'))) \quad (13)$$

In the special case of a uniform flake distribution $D(m) = 1/4\pi$, this model reduces to traditional “anisotropic” scattering, where the phase function depends only on the scattering angle: here, the scattering coefficient σ_s reduces to a constant related to the albedo

$$\sigma_s = a\rho \int_0^{\frac{\pi}{2}} \alpha(\cos \theta) \cos \theta \sin \theta d\theta,$$

the extinction coefficient turns into $\sigma_t = \frac{a\rho}{2}$, and f_p becomes a function of $\cos \theta = \omega \cdot -\omega'$

$$f_p(\cos \theta) = \frac{a\rho}{8\pi\sigma_s} \alpha\left(\cos \frac{\theta}{2}\right). \quad (14)$$

which can be solved for the α required to archive a particular f_p :

$$\alpha(\cos \theta) = \frac{8\pi\sigma_s}{a\rho} f_p(\cos(2\theta)) = 4\pi \frac{\sigma_s}{\sigma_t} f_p(\cos(2\theta)) \quad (15)$$

For instance, to represent the Henyey-Greenstein or Mie phase functions, equation (15) provides a way of encoding their properties into an albedo function α . Any isotropic medium can be expressed in this fashion, even when the material itself is not physically realizable




Figure 4: Two sample flake normal distributions and reflection profiles corresponding to the indicated incident direction ω_i : (a) $D(\omega) = \cos^{20}(v, \omega)_+$: Reflection similar to a rough surface, where v acts as surface normal. (b) $D(\omega) = \sin^{20}(v, \omega)$: Reflection similar to a rough fiber, where v acts as the fiber’s tangent direction. In both cases α is constant.

using flakes. For instance, by the previous equation, a medium without absorption and strong forward scattering will produce $\alpha > 1$ in some directions. Although the physical interpretation is lost in such cases, the framework remains internally consistent and all results hold regardless. This situation is similar to microfacet models, which can often still provide convenient abstractions when the underlying material breaks the statistical heightfield assumption.

Various types of anisotropic scattering can be modeled depending on the choice of flake distribution. For instance, concentrating $D(m)$ around a fixed direction (Figure 4a) produces behavior that resembles surface reflection, and in fact the form of the model is not unlike the commonly used Blinn-Phong BRDF, with the term $\sigma_t(\omega)$ playing the role of the usual cosine factor for surface illumination.

Similarly, it is also possible to produce fiber-like scattering by choosing a $D(m)$ that is concentrated near the plane perpendicular to the fiber’s tangent direction (Figure 4b). In this case, the resulting phase function resembles the reflection model of Kajiyama and Kay [1989], with $\sigma_t(\omega)$ playing the role of the usual sine factor for fiber illumination. Some differences exist; for instance, the microflake model produces a reflection cone that narrows towards forward scattering, resulting in less forward scattering than Kajiyama-Kay if $\alpha(\omega)$ is constant. However, the micro-flake approach does not share some of that model’s shortcomings, particularly its lack of energy conservation: especially for illumination near grazing angles, the Kajiyama-Kay model is far from normalized. Attempts to remedy this problem using a parameterized normalization constant invariably lead to a non-reciprocal phase function. The scattering model presented here, on the other hand, is always energy conserving, and it satisfies the appropriate reciprocity constraint (9) in conjunction with the previously defined directionally varying scattering coefficients.

Section 6 will present three solution techniques for rendering with this model, including a Monte Carlo approach and two approximate methods that involve the diffusion equation.

5 Anisotropic diffusion

To reason about diffusion in the presence of anisotropy, we must propagate the earlier changes to the RTE through the steps of this approximation. Classic derivations of the diffusion equation (e.g. Ishimaru [1978]) set the stage by considering solutions to the RTE under smoothness constraints.

The key assumption is that because each scattering event effectively blurs the incident illumination, high frequencies disappear from the angular radiance distribution as the number of scattering events increases. Motivated by this observation, the radiance function is restricted to a simple two-term expansion in terms of angular moments. Formally, for a function $f : S^2 \rightarrow \mathbb{R}$, the n -th moment on the unit sphere may be defined as

$$(\mu_n[f])_{i,j,k,\dots} := \int_{S^2} \underbrace{\omega_i \omega_j \omega_k \cdots}_{n \text{ factors}} f(\omega) d\omega. \quad (16)$$

For non-negative f , the 0th moment for instance gives the function's integral over the sphere, the 1st moment can be interpreted as a "center of mass" 3-vector, and the 2nd moment is a positive definite 3×3 -matrix. The above restriction then translates into setting

$$L(x, \omega) := \frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \omega \cdot E(x) \quad (17)$$

where the quantities $\phi(x) = \mu_0[L(x, \cdot)]$ and $E(x) = \mu_1[L(x, \cdot)]$ are known as *fluence* and *vector irradiance*, respectively.

The traditional derivation then proceeds by substituting this approximation into the RTE and enforcing equality amongst the 0th and 1st-order projections, which leads to the coupled set of equations

$$\nabla \cdot E(x) = Q_0(x) - \sigma_a \phi(x) \quad (18)$$

$$\nabla \phi(x) = 3Q_1(x) - 3\sigma'_t E(x) \quad (19)$$

where $Q_i(x) := \mu_i[Q(x, \cdot)]$, $\sigma'_t := \sigma_a + \sigma_s(1 - g)$ and the constant g specifies the *mean cosine* defined as $g := \int_{4\pi} f_p(\omega \cdot \omega')(\omega \cdot \omega') d\omega'$. Eliminating E in (19) produces the isotropic diffusion equation

$$\lambda \nabla^2 \phi(x) = \sigma_a \phi(x) - Q_0(x) + 3\lambda \nabla \cdot Q_1(x) \quad (20)$$

where $\lambda := \frac{1}{3\sigma'_t}$ is the diffusion constant.

5.1 Anisotropic diffusion—general form

The same set of ideas can also be applied to the anisotropic RTE. Here, the angular dependence of the scattering coefficients and the more general form of f_p causes the appearance of additional terms in the final set of equations. Since performing the two projections amounts to a mechanical application of rules also used in the isotropic derivation, we will only state the final result; the full derivation may be found in the supplementary material to this paper. The first and second-order projections of equation (2) are given by

$$\nabla \cdot E(x) = Q_0(x) - \frac{1}{4\pi} \mu_0[\sigma_a(x, \cdot)] \phi(x) - V(x) \cdot E(x) \quad (21)$$

$$\nabla \phi(x) = 3Q_1(x) - \frac{3}{4\pi} \mu_1[\sigma_a(x, \cdot)] \phi(x) - M(x) E(x) \quad (22)$$

where

$$V(x) := \frac{3}{4\pi} \left[\mu_1[\sigma_t(x, \cdot)] - \int_{S^2} \sigma_s(x, \omega) \mu_1[f_p(\cdot \rightarrow \omega)] d\omega \right]$$

$$M(x) := \frac{9}{4\pi} \left[\mu_2[\sigma_t(x, \cdot)] - \int_{S^2} \omega \sigma_s(x, \omega) \mu_1[f_p(\cdot \rightarrow \omega)]^T d\omega \right]$$

V is a 3-dimensional vector and M a 3×3 matrix, which both capture information about the low-frequency scattering behavior. In the case of isotropy, these equations together appropriately reduce to (20), and $M^{-1} = I/(3\sigma'_t)$ takes the role of the diffusion constant.

This approximation may be applied to any reflection model expressed in the framework of Section 3, and will be accurate to the extent that radiance does become uniform in the region of interest. In Section 5.3 we will further specialize to the micro-flake model.

5.2 Boundary conditions

We can proceed to derive an anisotropic generalization of the commonly used Robin boundary conditions: these specify that the net inward flux $\Gamma^{\text{in}}(x)$ matches the fraction F_{dr} of the outward flux that is reflected by interface effects $\Gamma^{\text{ref}}(x)$, i.e.

$$\Gamma^{\text{in}}(x) = \Gamma^{\text{ref}}(x) \quad (23)$$

where

$$\Gamma^{\text{in}}(x) = \int_{H^{2-}} L(x, \omega) |\omega \cdot n| d\omega \quad (24)$$

$$\Gamma^{\text{ref}}(x) = F_{dr} \int_{H^{2+}} L(x, \omega) |\omega \cdot n| d\omega \quad (25)$$

H^{2-} indicates integration over the inward hemisphere and the normal n is assumed to point outwards. Substituting the two-term expansion into (23) leads to

$$\frac{1}{4} \phi(x) - \frac{1}{2} n \cdot E(x) = F_{dr} \left[\frac{1}{4} \phi(x) + \frac{1}{2} n \cdot E(x) \right]$$

$$\Leftrightarrow \frac{1}{4} \phi(x) - \frac{A}{2} n \cdot E(x) = 0 \quad (26)$$

where we have defined $A := \frac{1+F_{dr}}{1-F_{dr}}$. For mismatched boundaries, we make use of the following rational approximation to measured data due to Egan and Hilgeman [1979]:

$$F_{dr} = -\frac{1.440}{\eta^2} + \frac{0.710}{\eta} + 0.668 + 0.0636\eta \quad (27)$$

where η is the relative refractive index at the boundary.

5.3 Anisotropic diffusion by micro-flakes

To apply the diffusion approximation to the micro-flake model, the coefficients in (21) and (22) must be computed. In this section, we derive these coefficients, discovering some important simplifications for micro-flakes compared to the general case.

Beginning with the coefficient of ϕ in (21), we find

$$\mu_0[\sigma_a(x, \cdot)] = 2a\rho\pi(1 - \hat{\alpha}) \quad (28)$$

where we have reversed the order of integration and defined

$$\hat{\alpha} := \int_0^\pi \alpha(|\cos \theta|) |\cos \theta| \sin \theta d\theta. \quad (29)$$

This constant may be interpreted as the overall reflection from a flake under constant illumination. When the albedo does not vary with respect to inclination, $\alpha = \hat{\alpha}$.

An immediate observation is that all first-order moments of the scattering and absorption coefficients vanish due to their mirror-symmetry (e.g. $\sigma_a(x, \omega) = \sigma_a(x, -\omega)$). This eliminates the first term in the expression for V . The remaining vector-valued integral

expression in $V(x)$ also vanishes: after a change of variables to half-directions, we have

$$V(x) = \frac{3a\rho}{4\pi} \int_{S^2} D(x, m) \int_{S^2} \omega' \alpha(|\omega' \cdot m|) |\omega' \cdot m| d\omega' dm = 0$$

which follows from the symmetry of the inner integrand (i.e. $f(\omega') = -f(-\omega')$).

By the same argument used for the first term of V , the coefficient of ϕ in (22) vanishes.

Turning finally to the matrix coefficient M , the first term in its definition takes on a simple form relating it to the low-frequency properties of $D(m)$:

$$\mu_2[\sigma_t(x, \cdot)] = \frac{a\rho\pi}{2} (I + \mu_2[D(x, \cdot)]). \quad (30)$$

The second term can be handled by a change of variables to half-directions and some simplifications, and then M is found to have a surprisingly simple closed form as a function of the second flake distribution moment:

$$M(x) = \frac{9a\rho}{8} ((1+2c_1+2c_2)\mu_2[D(x, \cdot)] + (1-2c_1)I). \quad (31)$$

where the constants

$$c_1 := \int_0^\pi \alpha(|\cos \theta|) |\cos \theta| \sin^3 \theta d\theta \quad (32)$$

$$c_2 := 2 \int_0^\pi \alpha(|\cos \theta|) |\cos \theta|^3 \sin \theta d\theta. \quad (33)$$

only depend on the choice of the albedo falloff curve. A property that will later be required is the positive definiteness of $M(x)$, which directly follows from $c_1, c_2 \geq 0$, $c_1 \leq \frac{1}{2}$ and the positive definiteness of $\mu_2[D(x, \cdot)]$.

When α is constant, $c_1 = \alpha/2$, $c_2 = \alpha$ and (31) further reduces to

$$M(x) = \frac{9a\rho}{8} ((1+3\alpha)\mu_2[D(x, \cdot)] + (1-\alpha)I). \quad (34)$$

The simplicity of (31) is of particular convenience when the flake distribution $D(m)$ is available in the form of a real spherical harmonics expansion: suppose that

$$D(x, \omega) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_l^m y_l^m(\omega)$$

where

$$y_l^m(\omega) = \begin{cases} Y_l^0(\omega), & \text{if } m = 0 \\ \sqrt{2} \operatorname{Re} Y_l^m(\omega), & \text{if } m > 0 \\ \sqrt{2} \operatorname{Im} Y_l^{-m}(\omega), & \text{if } m < 0 \end{cases}$$

and Y_l^m are the usual complex-valued orthonormal spherical harmonic basis functions. Then $\mu_2[D(x, \cdot)]$ is readily obtained in terms of the expansion coefficients:

$$\mu_2[D(x, \cdot)] = 2\sqrt{\frac{\pi}{15}} \begin{bmatrix} \sqrt{\frac{5}{3}} a_0^0 - \frac{a_2^0}{\sqrt{3}} + a_2^2 & a_2^{-2} & -a_2^{-1} \\ a_2^{-2} & \sqrt{\frac{5}{3}} a_0^0 - \frac{a_2^0}{\sqrt{3}} - a_2^2 & -a_2^{-1} \\ -a_2^{-1} & -a_2^{-1} & \sqrt{\frac{5}{3}} a_0^0 + 2\frac{a_2^0}{\sqrt{3}} \end{bmatrix}$$

Finally, with the $V(x)$ term gone, the equation also takes on a more familiar form: solving (22) for $E(x)$ and substituting into (21) as is done in the isotropic variant yields

$$\begin{aligned} \nabla \cdot [M(x)^{-1} (\nabla \phi(x) - 3Q_1(x))] \\ = \frac{1}{4\pi} \mu_0[\sigma_a(x, \cdot)] \phi(x) - Q_0(x). \end{aligned} \quad (35)$$

This step can be repeated once more to eliminate $E(x)$ from the boundary conditions (26):

$$\frac{1}{4} \phi(x) + \frac{A}{2} n^T M^{-1} (\nabla \phi(x) - 3Q_1(x)) = 0. \quad (36)$$

The two diffusion-based solution techniques in Section 6 are built on equation (35) and (36) and they only require that the problem be in this form. This means that as long as M is a positive definite matrix and $\mu_0[\sigma_a(x, \cdot)]$ a nonnegative constant, they can be applied irrespective of what model was used to generate them.

To summarize: we have derived a diffusion equation and suitable boundary conditions based on first- and second-order projections of our anisotropic RTE, which resulted in the introduction of an additional vector term V and the matrix M . For the micro-flake model, the V term reduced to zero and M took on a simple closed form solution (31) relating it to the low-frequency characteristics of the underlying flake distribution. With all of these pieces at hand, we are now able to compute approximate diffusion-based solutions to the anisotropic RTE when micro-flakes are used as the underlying phase function model.

6 Solution techniques

In this section, we present three different approaches for computing light transport in anisotropic media, spanning a range of performance and accuracy characteristics. First, we describe the necessary changes to support the flake model in a Monte Carlo renderer. Next, we adapt the finite element method to the anisotropic diffusion equation found in Section 5. And finally, we propose a new dipole BSSRDF for anisotropic translucent materials. Both diffusion approaches are general in the sense that they only depend on the cancellation of the $V(x)$ and $\mu_1[\sigma_a(x, \cdot)]$ terms in (21) and (22), which means that they may also be applied to any future scattering models sharing these properties.

6.1 Monte Carlo rendering

To use micro-flake distributions in a Monte Carlo rendering context, an efficient way of evaluating the coefficients σ_s , σ_t , and the phase function f_p will be necessary. Furthermore, good importance sampling support for f_p is crucial to achieve reasonably fast convergence, in particular when working with scatter-dominated media.

For simple families of flake distributions, it may be feasible to find closed-form importance sampling procedures and expressions for σ_s and σ_t . For general $D(m)$, however, we face a prohibitive numerical 2D integration for each scattering coefficient evaluation, since (11) and (12) are expressed in terms of spherical convolutions.

Both problems are greatly alleviated by moving to the frequency domain: for this, we represent $D(m)$ as a truncated spherical harmonics expansion using enough coefficients to ensure that the introduced error is negligible ($< 0.1\%$). The convolutions in the expressions for σ_s and σ_t are then found to involve the azimuthally invariant kernel functions $|\omega \cdot m|$ and $|\omega \cdot m| \alpha(\omega \cdot m)$, making it easy to find the coefficients of their respective expansions by an application of the Funk-Hecke theorem [Groemer 1996].

To support importance sampling of the phase function (13), we perform a 2D discretization over the outgoing direction: for each fixed outgoing ω , we tabulate the coefficients for a truncated spherical harmonics expansion of $f_p(\cdot \rightarrow \omega)$. During rendering, an expansion over the incident direction is obtained using bilinear interpolation from the nearest four outgoing direction samples. In all our examples we use between 7 and 10 coefficient bands and a 50×100 discretization over the $(\theta_\omega, \phi_\omega)$ -space.

With the spherical harmonics coefficients at hand, it is possible to apply the sampling technique proposed by Jarosz et al. [2009], which we found to work well in practice. Any procedure of this form of course introduces truncation and interpolation errors, and we therefore use the interpolated phase function only to guide the placement of samples, with (13) taken as the actual phase function during evaluation.

6.2 Finite Element Solution

The finite element method (FEM) provides a faster solution technique for the anisotropic diffusion equation. Our finite element solution has three components: first, we modify the Robin boundary condition (23) to produce a so-called anisotropic diffusive source boundary condition (ADSBC). Second, we use this new boundary condition to compute a matrix solution to the anisotropic diffusion equation (35). Finally, we present a query function that converts the computed solutions into radiance. Since our derivation closely follows the isotropic case [Arbree et al. 2009], we highlight only the differences here. More details can also be found in the supplementary material [Jakob et al. 2010].

6.2.1 Anisotropic Diffusive Source Boundary Condition

The ADSBC combines the basic Robin boundary condition with a new term describing illumination incident on the boundary due to light sources outside of the medium. This is often referred to as the *reduced intensity source*. This illumination is approximated by a diffusive flux term equal to the net refracted irradiance attenuated by the expected absorption before the first scattering event. This changes (23) to $\Gamma^{\text{in}}(x) = \Gamma^{\text{ref}}(x) + \Gamma_s(x)$ where

$$\Gamma_s(x) = e^{-\frac{\sigma_a(x,n)}{\sigma_s(x,n)}} \int_{H^{2+}} F_t(\omega) L(x, -\omega) (n \cdot \omega) d\omega \quad (37)$$

Following the mathematics presented in Section 5.2 ((24)–(36)) results in the anisotropic diffusive source boundary condition:

$$\frac{1}{4} \phi(x) + \frac{A}{2} n^T M^{-1} \nabla \phi(x) = \frac{1}{F_{dt}} \Gamma_s(x) \quad (38)$$

where $F_{dt} := 1 - F_{dr}$ denotes the overall Fresnel transmittance.

6.2.2 Anisotropic Finite Element Diffusion Equation

Next, this boundary condition is used to find the finite element solution of (22). This process has two steps: the derivation of the weak form and the discretization of that form into a matrix equation. Since these steps are essentially the same as in [Arbree et al. 2009], with (22) replacing the original isotropic equation, we present only the final matrix equation here:

$$(R + S + T)a = (q + g) \quad (39)$$

where

$$R_{ij} = \int_{\Omega} \nabla \beta_i \cdot (M^{-1} \nabla \beta_j) dx \quad q_i = \int_{\Omega} Q_0 \beta_i dx$$

and

$$S_{ij} = \frac{1}{4\pi} \int_{\Omega} \mu_0 [\sigma_a] \beta_i \beta_j dx \quad g_i = \frac{2}{AF_{dt}} \int_{\partial\Omega} \Gamma_s \beta_i ds$$

$$T_{ij} = \frac{1}{2A} \int_{\partial\Omega} \beta_i \beta_j ds$$

The β_i are basis functions defined at the vertices of the mesh and a denotes the solution's coefficient vector satisfying $\phi = \sum_{i=0}^{n-1} a_i \beta_i$. Since M^{-1} is always positive definite (Section 5.3), this equation is guaranteed to have a solution.

The final step is the derivation of the query function that converts the solution fluence of the matrix equation into the outgoing surface radiance (see [Jakob et al. 2010]). For completeness, we replicate this equation, which turns out to be the same as the isotropic case.

$$L(x, \omega) = \frac{F_t(\omega)}{4\pi} \left[\left(1 + \frac{1}{A}\right) \phi(x) - \frac{4}{F_{dt}A} \Gamma_s(x) \right] \quad (40)$$

Mathematically, this final solution is nearly identical to the isotropic case. The two key differences are the new definitions of the matrices R and S in (39).

6.3 BSSRDF model

In this section, we will first solve the anisotropic diffusion equation for a point source located in an infinite homogeneous medium. Following this, we use our solution to extend the BSSRDF dipole model proposed by Jensen et al. [2001] to anisotropic media.

6.3.1 Monopole

To arrive at the monopole solution, we model a point source located at the origin radiating the total amount of power Φ uniformly into all directions:

$$Q(x) = \frac{\Phi}{4\pi} \delta(x).$$

For this type of source, $Q_1(x) = 0$ which further simplifies the anisotropic diffusion equation (35). The mixed derivatives on the left-hand side of this equation can be eliminated by switching to a set of linearly transformed coordinates $\xi = Px$. The particular matrix P which accomplishes this is given by

$$P := QD^{\frac{1}{2}}Q^T \quad (41)$$

where QDQ^T is an eigendecomposition of M . This intuitively means that the anisotropic diffusion equation reduces to the isotropic variant when expressed using the appropriate coordinates:

$$\nabla^2 \phi(\xi) = \frac{1}{4\pi} \phi(\xi) \mu_0 [\sigma_a(\cdot)] - Q_0(\xi) \quad (42)$$

After a similar change of variables in the source term, its 0th-order projection is found to be $Q_0(\xi) = \Phi |\det P| \delta(\xi)$. Finally, the Green's function describing the fluence due to this type of source is found for example in Ishimaru [1978]:

$$\phi(x) = \frac{\Phi \det P}{4\pi \|Px\|} \exp\left(-\sqrt{\frac{1}{4\pi} \mu_0 [\sigma_a(\cdot)] \|Px\|}\right). \quad (43)$$

which is expressed back in terms of the original coordinates.




Figure 5: Illustration of the approximate boundary conditions used by the anisotropic BSSRDF model.

6.3.2 Dipole boundary conditions

We now switch to a half-space geometry with the interface located at $z = 0$ and a surface normal given by $n = (0, 0, 1)^T$. To describe this setup, the dipole model [Farrell and Patterson 1992] relies on an approximation to the Robin boundary conditions presented earlier: in the isotropic setting, Moulton [1990] has shown that (36) is approximately satisfied by requiring the fluence to be zero on an extrapolated boundary at distance $2A\lambda$ (where λ is the diffusion constant).

To find the corresponding boundary condition at the interface of an anisotropic medium, we can interpret (36) as a directional derivative of $\phi(x)$ along the direction $(M^{-1}n)/\|M^{-1}n\|$, i.e.

$$2A\|M^{-1}n\| \left(\frac{M^{-1}n}{\|M^{-1}n\|} \cdot \nabla \right) \phi(x) = -\phi(x).$$

When this approximation is good, we expect ϕ to vanish at a distance of $2A\|M^{-1}n\|$ in this direction. Since x was an arbitrary point on the surface, this property will in fact hold on a whole surface extrapolated at a distance of

$$d_p = 2A n^T M^{-1} n$$

which can be obtained by projecting the vector $M^{-1}n/\|M^{-1}n\|$ onto n (Figure 5).

6.3.3 Dipole solution

Following Farrell et al. [1992], an approximate half-space solution is now found based on a dipole source configuration. In this model, a negative ‘virtual’ light source is introduced to satisfy the zero fluence constraint on the extrapolated boundary.

In the isotropic case, placement of the sources is simple: the real source is positioned one mean free path below the surface ($z_r = -1/\sigma_t'$) to simulate the first scattering event due to a collimated beam at normal incidence. The zero fluence constraint at $z = 2A\lambda$ is satisfied by mirroring the real source along this plane, which leads to a virtual source position of $z_v = 1/\sigma_t' + 4A\lambda$.

The natural extension to anisotropic media places the real source at $z_r = -1/\sigma_t(n)$ below the surface, which matches the mean free path at normal incidence. Determination of the virtual source position is slightly more complicated since the monopole fluence (43) is now a function of $\|P(x - x_p)\|$. Under the linear transformation P , the extrapolated surface $z = d_p$ turns into a plane with a normal pointing into direction $p_1 \times p_2$, where p_i are the columns of P (Figure 6). Mirroring the light source along this plane and transforming the obtained position back into the original space results in

$$x_r := -\sigma_t(n)^{-1} n, \quad (44)$$

$$x_v := x_r + 2 \left(\sigma_t(n)^{-1} + d_p \right) \frac{P^{-1}(p_1 \times p_2)}{n^T (P^{-1}(p_1 \times p_2))} \quad (45)$$

The fluence due to both sources is then equal to

$$\phi(x) = \frac{\Phi \det P}{4\pi} \left[\frac{e^{-\beta d_r}}{d_r} - \frac{e^{-\beta d_v}}{d_v} \right] \quad (46)$$

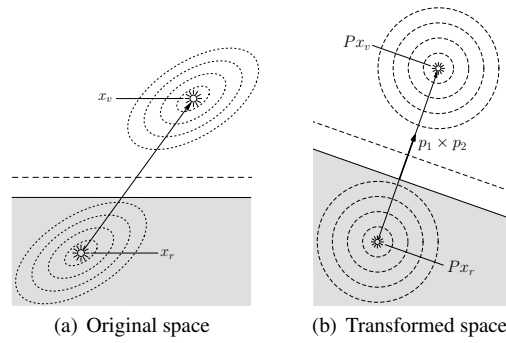


Figure 6: Geometric intuition of the anisotropic dipole: under a linear transformation, the diffusion becomes isotropic, which allows the straightforward construction of a dipole by mirroring the real source across the extrapolated surface (dashed line).

where $\beta := \sqrt{\frac{1}{4\pi} \mu_0 [\sigma_a(\cdot)]}$ and d_r and d_v denote distances computed in a space transformed by P , i.e. $d_r := \|P(x - x_r)\|$ and $d_v := \|P(x - x_v)\|$.

The final BSSRDF is obtained by computing the radiant exitance in the normal direction, $n \cdot E(x)$. Applying these steps leads to the familiar-looking

$$dM_o(x) = d\Phi \frac{\bar{\alpha}(n) \det P}{4\pi} \left[C_1 \frac{e^{-\beta d_r}}{d_r^3} + C_2 \frac{e^{-\beta d_v}}{d_v^3} \right], \quad (47)$$

where

$$C_1 := |n^T x_r| (\beta d_r + 1) \text{ and } C_2 := |n^T x_v| (\beta d_v + 1) \quad (48)$$

As in Jensen et al. [2001], we have added an albedo term $\bar{\alpha}(n) = \sigma_s(n)/\sigma_t(n)$ to account for the first scattering event in the earlier collimated beam approximation.

It is possible to further extend this model to multipole source configurations similar to Donner and Jensen [2005]. This is straightforward and entails simply adding more sources to cancel out the fluence on two extrapolated boundaries in the skewed solution space.

7 Results

Our final rendering pipeline, written in C++ and Java, ran on an 2.93Ghz Intel Xeon X5570 workstation with 8 cores. We implemented all three solution techniques, and this section presents an overview of our experimental results.

Scarf renderings: The scarf model was kindly provided by the authors of the recent work on yarn-level knit simulation [Kaldor et al. 2008]. To create volume densities based on their spline representation, we used a technique similar to the Lumislice algorithm [Xu et al. 2001], but took advantage of the flexibility of a general volume to build 2-ply yarns (yarns made by twisting multiple smaller yarns called ‘plies’). The fibers were well organized, so they should produce specular highlights running along the twisted plies. The final volume had an effective resolution of $1408 \times 256 \times 1152$ voxels, which required about 3 gigabytes of storage using a sparse grid data structure. Every occupied cell in the volume consisted of 4 single-precision values specifying density and local fiber orientation, and the latter was used both as an input to the Kajiy-Kay model, and to rotate flake distributions into the desired coordinate frame.

Three Monte Carlo simulations with full multiple scattering were performed at 512 samples per pixel using an isotropic phase function




Figure 7: (a), (b): Micro-flake and Kajiya-Kay renderings with exponents tuned for similar highlight sharpness in single-scattering images—already here, the lack of energy conservation in the Kajiya-Kay model is visible. (c): Kajiya-Kay rendering with full multiple scattering. (d): difference image showing the energy loss incurred in (c) as compared to the micro-flake rendering in Figure 1b.




Figure 8: (a), (b): Using spatially varying flake distributions oriented according to the direction maps captured by Marschner et al. [2005], we are able to reproduce the shifting reflection patterns (see the accompanying video) exhibited by their model, while furthermore accounting for multiple scattering. (c): Projecting collimated illumination onto the wood reveals interesting diffusion effects along the grain direction.

(Figure 1a), micro-flakes (Figure 1b) and the Kajiya-Kay model (Figure 7), respectively. The isotropic rendering lacks any specular effects, leading to dull and somewhat blurry appearance, while the anisotropic renderings show highlights delineating the plies. In the anisotropic case, multiple scattering naturally leads to near-white highlights (caused by single scattering), and in all cases it leads to increasing color saturation in the body of the material and then into the shadows, due to the many scattering events required for energy to reach these regions. The rendering using the Kajiya-Kay phase function shows similar local shading to the micro-flake model, since the two phase functions are similar. An important difference is that the micro-flake model is exactly normalized, whereas the Kajiya-Kay model loses energy towards grazing angles. A fraction of the total light energy is lost at every bounce, and this causes blurring in multiple scattering renderings, as well as a noticeable decrease in the overall image brightness (Figure 7d). As these examples demonstrate, accounting for oriented structures leads to a significant gain in visual realism, and the new model achieves this while retaining the full advantages of physics based rendering.

Changing from isotropic scattering to the micro-flake model when rendering heterogeneous materials does incur additional computational costs. Most of the extra time is spent evaluating $\sigma_t(\omega)$ at every step of ray marching between scattering events. In our code this currently involves evaluating a spherical harmonics expansion, leaving considerable opportunities for future optimization.

Wood example: As an example of a solid material with anisotropic structure, we used the measurements of finished wood released by Marschner et. al. [2005], which contain both diffuse

color and fiber direction textures for several types of wooden materials. To use these textures in a volume rendering context, the single scattering albedo α needs to be known, which we computed by numerically inverting the dipole BRDF approximation [Jensen et al. 2001] for each pixel of the diffuse color texture.

We then projected the curly maple textures into a slab filled with a flake medium and chose a distribution that simulates scattering from rough fibers ($D = \sin^{20} \theta$). In the resulting MC renderings (Figure 8), the flake medium is able to reproduce the shifting anisotropic reflections modeled by the original paper, and the ability to account for multiple scattering eliminates the need for an ad-hoc diffuse component. By illuminating the wood with a grid of collimated beam sources, we can observe anisotropic light diffusion occurring along the grain direction.

Finite element simulations: To validate the finite element computation against Monte Carlo, we used a simple scene to test the spreading of light from an illuminated spot. Because collimated illumination and half-spaces cannot be represented in the FEM system, we projected a 1mm diffuse spot on a 4×4 cm cube with material parameters $a\rho = 5.2482/cm$ and $\alpha = 0.998$. Except for the spot's immediate neighborhood, where the simulations are not expected to match because of the large contribution of low-order scattering, the solutions were found to be in good agreement (Figure 10a–c). Because of the structured sources in our examples, we modified the mesh refinement algorithm to increase resolution along edges in the projected source image. Since the finite element solution does not require any particular geometry or a homogeneous material, it is able to model a wide range of scenes, such as Figure 9a.




Figure 9: (a): Finite element solution to the Stanford Dragon illuminated by a grid of circular spots. The model is filled by an anisotropic micro-flake medium with its flake distribution aligned to a procedurally generated vector field, causing illumination to spread along curved paths. (b), (c): Comparison of the isotropic and anisotropic BSSRDF: a light source projects colored illumination on a translucent slab. Both use the same albedo and average density, but the anisotropic slab uses a \sin^{20} flake distribution about a diagonal axis parallel to the surface.

Scene	Algorithm	Rendering time
Scarf model (Kajiya-Kay)	MC	5 hr 24 min
Scarf model (Micro-flakes)	MC	22 hr 11 min
Heterogeneous dragon	FEM	6 min 26 sec
Isotropic slab (1 sample/pixel)	Dipole	3.4 sec
Anisotropic slab (1 sample/pixel)	Dipole	9.8 sec

Table 1: Timing results for some of our example scenes

Anisotropic dipole: We have added support for anisotropy to our existing implementation of the dipole BSSRDF [2001], which required only minimal changes to the code. Figure 9b–c shows the result of projecting colored illumination onto an anisotropic slab modeled in this way. In any dipole model renderer, a heuristic for transforming points into the half-space geometry is required, and in the anisotropic case a simple solution based on distance is insufficient, since dM_o is now a function of two parameters instead of just the radial distance. To keep our implementation simple, we simply drop one coordinate, and this can be seen to introduce artifacts on the sides of the slab. We believe that a heuristic with overall smoother behavior could be used to remove such artifacts.

Similarly to the FEM test case, we constructed a simple scene to validate the dipole model, and we ran Monte Carlo simulations of a collimated beam impinging on an anisotropic half-space with the same parameters as before. Figure 10d–f shows a result of this comparison; as in the isotropic case, the dipole is only found to be accurate in the far field. A curious observation in our simulations was that the lower-order scattering caused by the very concentrated beam source tends to propagate perpendicularly to the main direction of diffusion (Figure 10e), which was also noticed by Kienle et al. [2004]. The dipole does not attempt to model this effect, which accounts for some of the differences. We do not see this as an obstacle for graphics applications, since the lower-order scattering effect quickly disappears as one moves to finite beam widths.

Another noteworthy difference is that the anisotropic BSSRDF may not achieve its maximum value at the point of incidence, which happens when increased diffusion occurs along an axis that is tilted with respect to the surface plane. This is not necessarily a failure of the dipole, but a materialization of the disagreement between the embedded point source and collimated beam solutions in such media. When this effect is not desired, the dipole solution can be re-centered to remove the inconsistency. A more extensive overview of the Monte Carlo comparisons can be found in the accompanying technical report [Jakob et al. 2010].

8 Conclusion

This paper has provided the theory and practice required to integrate anisotropic media into physically based rendering systems. An important piece of this is the introduction of a new model for media, based on micro-flake distributions, which generalizes previous phase function models and provides a physically sound alternative to the more heuristic reflectance models that are normally used in volume models of materials with directional scattering properties.

We also have shown how the more general anisotropic radiative transfer equation still can be approximated by diffusion (though a diffusion matrix is required where a scalar diffusion constant was formerly used), which provides fast, approximate solutions to multiple scattering in highly scattering media, by using a finite element procedure or by using a more approximate but extremely fast anisotropic dipole solution. The diffusion approximation shares the limitations familiar from the isotropic case: it applies only when radiance becomes uniform in direction, which happens only in high albedo media and only far from the source—and these are satisfied more quickly the lower the degree of anisotropy.

The framework introduced in this paper provides a basis for future work in several areas. First of all, new volume models for complex geometry at any scale, from cloth viewed at arm’s length to a forest viewed from across a valley, can now be built based on physically realistic scattering models, with multiple scattering naturally adding richness to the appearance, rather than being built as volume aggregations of surfaces shaded by heuristic single-scattering models. Volume models are also a natural setting for multiresolution approaches [Neyret 1998]. Aggregating geometry to larger scales can lead to anisotropic behavior even when the volume is isotropic at the small scale, so our anisotropic framework provides an ideal platform for building volume level-of-detail in a physically based setting.

The framework we have presented does have limitations. One is computational cost: volume models consume significant memory, and the more so when they are augmented with spatially varying vectors or other parameters for anisotropy; also, evaluating anisotropic models consumes more computing time than looking up scalar parameters. The radiative transfer formulation itself has limitations, and may not correctly model short-distance effects below the correlation length of the random medium being simulated. The diffusion approximation has further limitations; it only can be expected to be accurate when the radiance distribution becomes directionally smooth, which happens less easily with highly anisotropic or sharply heterogeneous media, and the dipole solution, while it gives reasonable estimates in the far field for half-space geometry, is (just as in

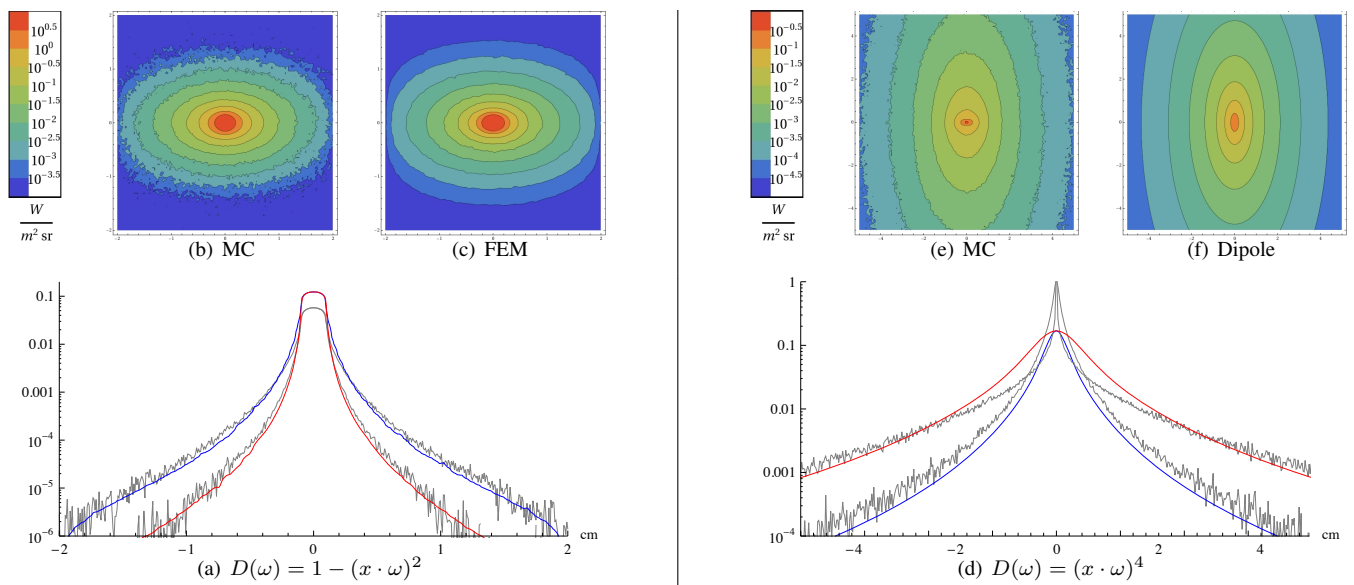


Figure 10: Validation of the finite element and dipole solutions against Monte Carlo references for two different distributions $D(\omega)$ centered around the x -axis. The blue and red graphs represent horizontal and vertical slices through the origin, respectively. (a)–(c): The finite element diffusion solution agrees well with the Monte Carlo simulation away from the source. (d)–(f): The dipole solution is more approximate but provides a useful match far from the source. The Monte Carlo simulation shows a characteristic feature of anisotropic media: low-order scattering produces horizontal spread near the center and high-order scattering produces vertical spread in the diffusion region.

the isotropic case) only a very rough model when applied to more complex surfaces.

Also, though the micro-flake model has considerable flexibility and can capture the basic phenomena of surface-like, fiber-like, and point-like scattering, not all volume scattering functions can be represented by micro-flakes. Future work will need to discover the practical significance of its limits in representing realistic materials, developing extensions where needed. Aggregating given geometry into volume models (another form of volume level of detail) by selecting appropriate micro-flake distributions is also an important potential application.

The most important implication of our work for the future of volume models is that they can now act as full-fledged participants in global illumination systems. This means volume models, which have proven so successful for fur, cloth, trees, and other complex structures, can take full advantage of the richness and realism that is naturally provided by complete light transport simulations.

9 Acknowledgements

This work was supported by the National Science Foundation (grants CCF-0347303 and CCF-0541105) and by Unilever Corporation. The authors also thank Jonathan Kaldor, Manuel Vargas and Manolis Savva for providing the scarf dataset.

References

- ARBREE, A., WALTER, B., AND BALA, K. 2009. Heterogeneous subsurface scattering using the finite element method. To appear in *IEEE Transactions on Visualization and Computer Graphics*.
- ARRIDGE, S. R. 1999. Optical tomography in medical imaging. *Inverse Problems* 15, 2, R41–R93.

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000*, 65–74.

- BLINN, J. F. 1982. Light reflection functions for simulation of clouds and dusty surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 82)*, 21–29.

- CEREZO, E., PÉREZ, F., PUEYO, X., SERON, F. J., AND SILLION, F. X. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5, 303–328.

- CHANDRASEKHAR, S. 1959. *Radiative Transfer*. Oxford University Press.

- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (Jan.), 7–24.

- CRASSIN, C., NEYRET, F., LEFEBVRE, S., AND EISEMANN, E. 2009. Gigavoxels : Ray-guided streaming for efficient and detailed voxel rendering. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D)*.

- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 1032–1039.

- DREBIN, R. A., CARPENTER, L., AND HANRAHAN, P. 1988. Volume rendering. In *Computer Graphics (Proceedings of SIGGRAPH 88)*, 65–74.

- DUDKO, O. K., AND WEISS, G. H. 2005. Estimation of anisotropic optical parameters of tissue in a slab geometry. *Biophysical Journal* 88, 5, 3205 – 3211.

- EGAN, W., AND HILGEMAN, T. 1979. *Optical Properties of Inhomogeneous Materials*. Academic Press New York.

- FARRELL, T. J., AND PATTERSON, M. S. 1992. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for

- the noninvasive determination of tissue optical properties in vivo. In *Medical Physics, Volume 19, Issue 4*, 879–888.
- FEDKIW, R., STAM, J., AND JENSEN, H. W. 2001. Visual simulation of smoke. In *Proceedings of ACM SIGGRAPH 2001*, 15–22.
- GIBSON, A. P., HEBDEN, J. C., AND ARRIDGE, S. R. 2005. Recent advances in diffuse optical imaging. *Physics in Medicine and Biology* 50, R1–R43.
- GROEMER, H. 1996. *Geometric applications of Fourier series and spherical harmonics*. Cambridge Univ Press.
- HEINO, J., ARRIDGE, S., SIKORA, J., AND SOMERSALO, E. 2003. Anisotropic effects in highly scattering media. *Phys. Rev. E* 68, 3.
- HEISKALA, J., NISSILA, I., NEUVONEN, T., JARVENPAA, S., AND SOMERSALO, E. 2005. Modeling anisotropic light propagation in a realistic model of the human head. *Applied Optics* 44, 11, 2049–2057.
- ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media*. Academic Press, New York, USA.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. Tech. rep., Cornell University. (Expanded version), <http://hdl.handle.net/1813/14982>.
- JAROSZ, W., CARR, N. A., AND JENSEN, H. W. 2009. Importance Sampling Spherical Harmonics. *Computer Graphics Forum (Proc. Eurographics EG'09)* 28, 2 (4), 577–586.
- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of 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 SIGGRAPH 01*, 511–518.
- JOHNSON, P. M., AND LAGENDIJK, A. 2009. Optical anisotropic diffusion: new model systems and theoretical modeling. *Journal of Biomedical Optics* 14, 5.
- KAJIYA, J. T., AND HERZEN, B. P. V. 1984. Ray tracing volume densities. In *Computer Graphics (Proceedings of SIGGRAPH 84)*, 165–174.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 271–280.
- KALDOR, J. M., JAMES, D. L., AND MARSCHNER, S. 2008. Simulating knitted cloth at the yarn level. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, 1–9.
- KIENLE, A., FORSTER, F., AND HIBST, R. 2004. Anisotropy of light propagation in biological tissue. *Optics letters* 29, 22, 2617–2619.
- LACROUTE, P., AND LEVOY, M. 1994. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proceedings of SIGGRAPH 94*, 451–458.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Eurographics Rendering Workshop 1996*, 91–100.
- LEVOY, M. 1988. Display of surfaces from volume data. *IEEE Computer Graphics & Applications* 8, 3 (May), 29–37.
- LI, P., AND UREN, N. F. 1998. Analytical solution for the electric potential due to a point source in an arbitrarily anisotropic half-space. *Journal of Engineering Mathematics* 33, 129–140.
- MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph.* 24 (July), 727–734.
- MAX, N. L. 1994. Efficient light propagation for multiple anisotropic volume scattering. In *Fifth Eurographics Workshop on Rendering*, 87–104.
- MISHCHENKO, M. I., TRAVIS, L. D., AND LACIS, A. A. 2006. *Multiple Scattering of Light by Particles*. Cambridge U. Press.
- MOULTON, J. 1990. *Diffusion modeling of picosecond laser pulse propagation in turbid media*. Master's thesis, McMaster University.
- NEYRET, F. 1998. Modeling, animating, and rendering complex scenes using volumetric textures. *IEEE Transactions on Visualization and Computer Graphics* 4, 1 (Jan./Mar.), 55–70.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, 11–22.
- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 253–262.
- PREISENDORFER, R. 1976. *Hydrologic Optics*. US Dept Commerce.
- PREMOZE, S., ASHIKHMIN, M., TESSENDORF, J., RAMAMOORTHY, R., AND NAYAR, S. 2004. Practical rendering of multiple scattering effects in participating media. In *Rendering Techniques 2004: 15th Eurographics Workshop on Rendering*, 363–374.
- RUSHMEIER, H. E., AND TORRANCE, K. E. 1987. The zonal method for calculating light intensities in the presence of a participating medium. In *Computer Graphics (Proceedings of SIGGRAPH 87)*, 293–302.
- RYZHIK, L., PAPANICOLAOU, G., AND KELLER, J. B. 1996. Transport equations for elastic and other waves in random media. *Wave Motion* 24, 1–44.
- SUN, B., RAMAMOORTHY, R., NARASIMHAN, S. G., AND NAYAR, S. K. 2005. A practical analytic single scattering model for real time rendering. *ACM Trans. Graph.* 24, 3 (Aug.), 1040–1049.
- VAN DE HULST, H. C. 1957. *Light Scattering by Small Particles*. John Wiley & Sons.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Eurographics Workshop on Rendering 2007*, 195–206.
- WESTOVER, L. 1990. Footprint evaluation for volume rendering. In *Computer Graphics (Proceedings of SIGGRAPH 90)*, 367–376.
- XU, Y.-Q., CHEN, Y., LIN, S., ZHONG, H., WU, E., GUO, B., AND SHUM, H.-Y. 2001. Photo-realistic rendering of knitwear using the lumislice. In *Proc. ACM SIGGRAPH 2001*, 391–398.