

Predicting Reflectance Functions from Complex Surfaces

Stephen H. Westin
James R. Arvo
Kenneth E. Torrance

Program of Computer Graphics
Cornell University
Ithaca, New York 14853

Abstract

We describe a physically-based Monte Carlo technique for approximating bidirectional reflectance distribution functions (BRDFs) for a large class of geometries by directly simulating optical scattering. The technique is more general than previous analytical models: it removes most restrictions on surface microgeometry. Three main points are described: a new representation of the BRDF, a Monte Carlo technique to estimate the coefficients of the representation, and the means of creating a milliscale BRDF from microscale scattering events. These allow the prediction of scattering from essentially arbitrary roughness geometries. The BRDF is concisely represented by a matrix of spherical harmonic coefficients; the matrix is directly estimated from a geometric optics simulation, enforcing exact reciprocity. The method applies to roughness scales that are large with respect to the wavelength of light and small with respect to the spatial density at which the BRDF is sampled across the surface; examples include brushed metal and textiles. The method is validated by comparing with an existing scattering model and sample images are generated with a physically-based global illumination algorithm.

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

Additional Key Words: spherical harmonics, Monte Carlo, anisotropic reflection, BRDF

1 Introduction

Since the earliest days of computer graphics, experimenters have recognized that the realism of an image is limited by the sophistication of the model of local light scattering [3, 12]. Non-physically-based local lighting models, such as that of Phong [12], although computationally simple, exclude many important physical effects and lack the energy consistency needed for global illumination calculations. Physically-based models [2, 5, 15] reproduce many effects better, but cannot




Figure 1: Applicability of Techniques

model many surfaces, such as those with anisotropic roughness. Models that deal with anisotropic surfaces [8, 11] fail to assure physical consistency.

This paper presents a new method of creating local scattering models. The method has three main components: a concise, general representation of the BRDF, a technique to estimate the coefficients of the representation, and a means of using scattering at one scale to create a BRDF for a larger scale. The representation used makes it easy to enforce the basic physical property of scattering reciprocity, and its approximation does not require discretizing scattering directions as in the work of Kajiyama [8] and Cabral et al. [1].

The method can predict scattering from any geometry that can be ray-traced: polygons, spheres, parametric patches, and even volume densities. Previous numerical techniques were limited to height fields, and analytical methods have been developed only for specific classes of surface geometry. The new method accurately models both isotropic and anisotropic surfaces such as brushed metals, velvet, and woven textiles.

Figure 1 shows several representations used in realistic rendering, along with approximate scale ranges where each is applicable. At the smallest scale (size $\ll 1$ mm), which we call *microscale*, the BRDF accurately captures the appearance of a

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

surface. As individual surface features become larger than one pixel, texture maps, bump maps, and texels can be used to show surface features. At the largest scale, *object scale*, the geometry must be modeled explicitly, for example with polygons or parametric patches.

The applicability of each representation ultimately depends on the context: the upper limit of applicable scale is determined by the frequency of sampling across a surface, and the lower limit is determined by the integration area for each sample; this is often the surface area represented by a pixel. When rendering, say, an interior scene, objects as small as a pencil must be modeled at object scale; when simulating the view from orbit, however, objects as large as trees and buildings can be modeled within the BRDF, so we can think of them as microscale geometry, or *microgeometry*. The advent of global illumination methods (e.g. [6, 18]) has created another concept of scale: these methods generally use a coarser characterization of scattering for indirect illumination, but demand careful attention to energy consistency and physical accuracy.

The method of this paper is applicable wherever the BRDF is an adequate model of surface geometry. It uses an analytical BRDF model for scattering at one scale of roughness, the *microscale*, simulating geometric optical scattering at a larger scale, the *milliscale*. Milliscale scattering embodies large-scale roughness effects (roughness size \gg wavelength of light, λ), and any smooth surface effects (roughness size $\approx \lambda$ or $< \lambda$) are modeled by the microscale BRDF, which can include wave optics effects.

The next three sections present the heart of the technique: the BRDF representation, the Monte Carlo estimator, and the means of estimating a milliscale BRDF from the microscale description of surface roughness.

2 Wheels Within Wheels: Representing the BRDF with Spherical Harmonics

A general scattering function for unpolarized light is a function of four variables, $\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) : S^2 \times S^2 \mapsto \mathcal{R}$, where S^2 is the unit sphere, θ_i, ϕ_i are the elevation and azimuth angles of incidence, and θ_r, ϕ_r are the corresponding angles of reflection (Figure 2). For a BRDF, ρ_{bd} is zero whenever θ_i or $\theta_r > \frac{\pi}{2}$. The BRDF can take on highly arbitrary shapes [5, 16], so a very general method is needed to represent it. Fortunately, a BRDF




Figure 2: Scattering Angles

is generally smooth, making it a good candidate for representation by smooth orthogonal functions. Previous authors have used spherical harmonics to represent scattering functions [1, 9, 13], since they form a complete basis set of smooth functions over the sphere. Kajiya [9] used spherical harmonics to derive an analytical scattering function; Cabral et al. [1] and Sillion et al. [13] used them as a numerical approximation to the BRDF. The representation used in this paper is an extension of Sillion's technique; it provides an accurate, concise embodiment of the general BRDF.

2.1 Overview of Spherical Harmonics

Any square-integrable function over the sphere can be exactly represented by an infinite sum of spherical harmonic basis functions, $Y_{lm}(\theta, \phi)$, of varying order, l , and degree, m :

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l C_{lm} Y_{lm}(\theta, \phi). \quad (1)$$

As with a Fourier representation, we can approximate f by truncating the series to a finite number of terms. For convenience, we organize this finite collection of basis functions into a vector by the convention of encoding both order and degree with a single subscript. Thus

$$f(\theta, \phi) \approx \sum_{k=0}^n C_k Y_k(\theta, \phi) = \mathbf{C} \cdot \mathbf{Y}(\theta, \phi). \quad (2)$$

Each coefficient C_k is defined by the inner product of $f(\theta, \phi)$ with the corresponding spherical harmonic basis function:

$$\begin{aligned} C_k &= \int_0^{2\pi} \int_0^{\pi} f(\theta, \phi) Y_k(\theta, \phi) \sin \theta \, d\theta \, d\phi \\ &= \langle Y_k | f \rangle. \end{aligned} \quad (3)$$

This follows directly from the orthogonality of the basis functions [17].

2.2 Representing the BRDF

If we fix the incident direction (θ_i, ϕ_i) , the BRDF is a function of two variables, (θ_r, ϕ_r) , and the representation in Equation 2 suffices. To account for variation of the BRDF with incident direction, the coefficient vector \mathbf{C} in Equation 2 can be thought of as a function of the incident direction. If a surface has isotropic roughness, as assumed in [1] and [13], the scattering function ρ_{bd} is independent of rotation about the surface normal. In this case,

$$\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) = \rho_{bd}(\theta_i, 0, \theta_r, \phi_r - \phi_i). \quad (4)$$

Each coefficient C_k is a function of θ_i alone, which can be calculated for a number of selected values of θ_i and interpolated for all θ_i [1, 13]. In general, however, a BRDF is a function of ϕ_i as well as of $(\theta_i, \theta_r, \phi_r)$, so a richer representation is needed.

2.3 Extension to Anisotropic Surfaces

For an anisotropic surface both θ_i and ϕ_i must be considered, and none of the previous representations suffices [1, 13]. Each coefficient C_k in Equation 2 is thus a function of θ_i and ϕ_i :

Nomenclature

C_k	Spherical harmonics coefficient for basis function Y_k
\mathbf{C}	Vector of coefficients
E_i	Incident energy flux density (irradiance)
E_r	Reflected energy flux density
G_k	Estimator of C_k
I_i	Incident radiance
I_r	Reflected radiance
\mathbf{M}	Exact matrix of coefficients to represent ρ_{bd}
$\widetilde{\mathbf{M}}$	Monte Carlo approximation of \mathbf{M}
m_{jk}	Element at row j , column k of matrix \mathbf{M}
N_b	Number of exit rays resulting from one incident ray
N_i	Number of incident ray directions
N_p	Number of sample points on surface of specimen
$p(\theta_i, \phi_i, \theta_r, \phi_r)$	Probability density function of scattering from (θ_i, ϕ_i) to (θ_r, ϕ_r)
$R(\theta_i, \phi_i, \theta_r, \phi_r)$	Attenuation of a single ray incident from (θ_i, ϕ_i) and reflected to (θ_r, ϕ_r)
$Y_k(\theta, \phi)$	Spherical harmonics basis function
\mathbf{Y}	Vector of basis functions
θ	Elevation angle: $\theta = 0$ at surface normal
ϕ	Azimuth angle: $\phi = 0$ at x axis
$\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r)$	Milliscale bidirectional reflectance distribution function (BRDF)
$\hat{\rho}_{bd}(\theta_i, \phi_i, \theta_r, \phi_r)$	Microscale bidirectional reflectance distribution function (BRDF)
$\hat{\rho}_s(\theta_i, \phi_i)$	Microscale specular reflectivity
$d\omega_i$	Differential solid angle of incident energy
$d\omega_r$	Differential solid angle of reflected energy
$\langle a b \rangle$	Inner product of two functions: $\int a(t)b(t)dt$
$\langle \xi \rangle$	Expected value of random variable ξ

$$\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) \approx \sum_{k=0}^n C_k(\theta_i, \phi_i) Y_k(\theta_r, \phi_r). \quad (5)$$

Each coefficient function, $C_k(\theta_i, \phi_i)$, is defined by the inner product of $\rho_{bd}(\theta_i, \phi_i, \cdot, \cdot)$ with the corresponding spherical harmonic basis function:

$$C_k(\theta_i, \phi_i) = \langle \rho_{bd} | Y_k \rangle_{refl} \quad (6)$$

where the subscript “*refl*” denotes integration over the reflected hemisphere. Reciprocity makes the dependence of ρ_{bd} on (θ_i, ϕ_i) exactly like its dependence on (θ_r, ϕ_r) . Since spherical harmonics concisely represent the latter dependence, we also use them to represent the dependence on (θ_i, ϕ_i) , expressing each coefficient function in terms of spherical harmonics. Each element of our vector \mathbf{C} of coefficients is now represented in turn by a vector of coefficients, giving us a matrix \mathbf{M} to represent the BRDF. Each element of the matrix \mathbf{M} is given by

$$m_{jk} = \langle Y_j | \langle \rho_{bd} | Y_k \rangle_{refl} \rangle_{in} \quad (7)$$

where the subscripts “*in*” and “*refl*” denote integration over the incident and reflected hemispheres, respectively. Evaluation of the BRDF becomes

$$\begin{aligned} \rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) &\approx \sum_{j=0}^N \sum_{k=0}^N Y_j(\theta_i, \phi_i) m_{jk} Y_k(\theta_r, \phi_r) \\ &= \mathbf{Y}^T(\theta_i, \phi_i) \mathbf{M} \mathbf{Y}(\theta_r, \phi_r), \end{aligned} \quad (8)$$

where $\mathbf{Y}(\theta, \phi)$ is the column vector of basis functions evaluated at (θ, ϕ) .

2.4 Reciprocity

An important physical constraint on the BRDF is *reciprocity*, which states that

$$\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) = \rho_{bd}(\theta_r, \phi_r, \theta_i, \phi_i) \quad (9)$$

for all angles of incidence and reflection [14]. If the matrix \mathbf{M} is symmetric, then

$$\mathbf{Y}^T(\theta_i, \phi_i) \mathbf{M} \mathbf{Y}(\theta_r, \phi_r) = \mathbf{Y}^T(\theta_r, \phi_r) \mathbf{M} \mathbf{Y}(\theta_i, \phi_i) \quad (10)$$

and the approximation in Equation 8 satisfies Equation 9. By assuring that we compute a symmetric matrix \mathbf{M} , we can enforce exact reciprocity; previous approaches [1, 8, 11, 13] afforded, at best, approximate reciprocity.

2.5 Storage Reduction and Filtering

The matrix \mathbf{M} can be quite large; tens of thousands of elements are typical. Since our BRDF representation, like that of [13], is based on spherical harmonics, we can adapt two techniques from that work to reduce the number of coefficients

(and corresponding basis functions) needed: the first technique causes half the coefficients to vanish, and the second reduces the high-frequency content of the BRDF, reducing the number of coefficients needed to achieve an acceptably accurate approximation. Since we deal only with scattering to one hemisphere, we can complete the other hemisphere with an arbitrary function. We chose a function that reduces the size of the representation: $\rho_{bd}(\theta_i, \phi_i, \pi - \theta_r, \phi_r) = -\rho_{bd}(\theta_i, \phi_i, \theta_r, \phi_r)$; this causes half of the coefficients (those with $l + m$ even in the real form of spherical harmonics) to be zero; they can be omitted from the representation, reducing the matrix size by $\frac{3}{4}$. To economize further, we represent $\rho_{bd} \cos \theta_i \cos \theta_r$ instead of ρ_{bd} ; multiplication by $\cos \theta_i$, together with the completion described above, forces C^1 continuity at the equator and drastically reduces ringing. To maintain symmetry of the matrix \mathbf{M} , we also multiply by $\cos \theta_r$. Representing $\rho_{bd} \cos \theta_i \cos \theta_r$ assures that Equation 9 is still satisfied. We omit this implementation detail from the following discussion.

As with a Fourier representation of a function, simply truncating all coefficients with index $l > l_{max}$ will cause ringing in the approximation, called the Gibbs phenomenon. To reduce this, we attenuate higher frequencies, as did Cabral et al. [1], by progressively reducing the magnitude of coefficients with $l_{filter} < l \leq l_{max}$, where l_{filter} is an empirically-determined threshold. The magnitude is reduced according to a half-Gaussian with empirically-determined width.

3 Monte Carlo Estimation of the Coefficient Matrix

If we bombard a specimen with incident rays from an arbitrary direction $U = (\theta_i, \phi_i)$, the BRDF can be expressed as

$$\rho_{bd}(U, V) = \frac{p(U, V) \langle R(U, V) \rangle}{\cos \theta_r} \quad (11)$$

where a ray from direction U will scatter into $V = (\theta_r, \phi_r)$ with a probability density $p(U, V)$, and $\langle R(U, V) \rangle$ is the mean attenuation of all rays incident from direction U and scattered in direction V .

In order to obtain a spherical harmonics coefficient, we must integrate the product $\rho_{bd} Y_k$ over the hemisphere.

$$\begin{aligned} C_k(U) &= \int_{S^2} \rho_{bd}(U, V) Y_k(V) dV \\ &= \int_{S^2} \langle g_k(U, V) \rangle p(U, V) dV \end{aligned} \quad (12)$$

where

$$g_k(U, V) = \frac{R(U, V)}{\cos \theta_r} Y_k(V). \quad (13)$$

Unfortunately we have no analytical expression for p or R ; we can, however, use a Monte Carlo simulation to estimate the integral in Equation 12. The integral can be interpreted as the expected value of $g_k(U, V)$, where V is a random variable with probability density function $p(U, V)$. If we define

$$G_k(U) = \frac{1}{N} \sum_{n=1}^N g_k(U, V_n) \quad (14)$$

where V_n are random samples distributed according to p , then the expected value of G_k is $C_k(U)$; G_k is said to be an *estimator*

of the integral [10]. The rays departing from the specimen in direction V will have mean attenuation $\langle R(U, V) \rangle$; this attenuation must be multiplied by $Y_k(V) / \cos \theta_r$ to give the expected value g for the estimator.

This leaves another integration, that with respect to U :

$$m_{jk} = \int_{S^2} C_k(U) Y_j(U) dU. \quad (15)$$

This integration can also be handled via Monte Carlo, this time as *quadrature*, a discrete approximation to an integral. This is handled similarly, with the estimator

$$\tilde{m}_{jk} = \frac{1}{N} \sum_{n=1}^N C_j(U_n) Y_j(U_n) \quad (16)$$

where the U_n are uniformly distributed over the incident hemisphere. These two sampling processes, each approximating an integral in two dimensions, can be combined into one process to approximate the four-dimensional integral desired.

$$\tilde{m}_{jk} = \frac{1}{N} \sum_{n=1}^N g_k(U_n, V_n) Y_j(U_n) \quad (17)$$

where the U_n are distributed uniformly and the V_n are distributed according to p . \tilde{m}_{jk} is an unbiased estimator of m_{jk} .

The simulation yields $\tilde{\mathbf{M}}$, an approximation to the symmetric matrix \mathbf{M} , and does not guarantee symmetry, so reciprocity of the BRDF is not guaranteed. We average the upper triangle and the lower triangle of $\tilde{\mathbf{M}}$ to obtain a symmetric matrix $\frac{1}{2}(\tilde{\mathbf{M}} + \tilde{\mathbf{M}}^T)$ which is used to compute ρ_{bd} . The two triangles are independent unbiased estimates of the BRDF; by averaging them to obtain a symmetric matrix, we also reduce the variance of our estimate of \mathbf{M} .

4 From Microscale to Milliscale

The BRDF can be used to model features ranging from microscale to milliscale for visible light, as shown in Figure 1. This section explains how to use microscale scattering events to calculate a milliscale BRDF. The section starts with basic BRDF definitions, describes the individual microscale scattering events, then explains how individual Monte Carlo events are incorporated into the milliscale model to obtain an aggregate BRDF.

At the microscale, arbitrary reflection models may be employed, including ideal specular, ideal diffuse, and directional diffuse models. One illustrative case is where the microgeometry is composed of planar ideal specular surfaces; this is equivalent to geometric optics models based on microfacets, such as the Torrance-Sparrow model [15].

We use ray tracing to model scattering events, as suggested by Cabral et al. [1]. The ray tracer must be carefully designed to assure physically accurate results. Each ray has a certain amount of energy associated with it; microscale reflection will attenuate this energy and perhaps divide it among multiple rays at each bounce. All calculations involve energy flux density until a ray finally exits the model; then the energy is converted to radiance, the proper quantity for the BRDF, by dividing by $\cos \theta_r$. The radiance distribution is averaged over the specimen surface to create a milliscale BRDF.




Figure 3: Local Scattering Modes

4.1 Incident Energy and the BRDF

We are estimating the BRDF ρ_{bd} , which is expressed at a given wavelength as

$$\rho_{bd}(\theta_r, \phi_r, \theta_i, \phi_i) = \frac{dI_r(\theta_r, \phi_r)}{dE_i(\theta_i, \phi_i)} \quad (18)$$

where dI_r is the reflected radiance and dE_i is the incident energy flux density, the incident energy per unit time per unit area. This equation holds at both micro- and milliscales. It becomes simpler to evaluate if we hold the denominator (incident energy flux density) constant and vary the incident angles θ_i, ϕ_i . Then

$$\rho_{bd}(\theta_r, \phi_r, \theta_i, \phi_i) = \frac{dI_r(\theta_r, \phi_r)}{dE_i} \quad (19)$$

where dE_i is the (constant) incident energy flux density.

Incident radiance I_i is defined as the incident energy flux density per unit projected area per unit solid angle

$$I_i = \frac{dE_i}{\cos \theta_i d\omega_i} \quad (20)$$

Thus

$$dE_i(\theta_i, \phi_i) = I_i(\theta_i, \phi_i) \cos \theta_i d\omega_i \quad (21)$$

The factor $\cos \theta_i$ converts receiving area to projected area, accounting for the dependence of projected surface area on θ_i .

The method allows different local scattering modes, three of which are shown in Figure 3. The next three sections describe how these modes are modeled.

4.2 Specular Reflection

The BRDF at the microscale may contain an ideal specular component $\hat{\rho}_s$. Whenever a ray hits such a microfacet, we model the transfer by spawning a ray in the specular direction as in classical ray-tracing [19]. The energy flux density of this ray is determined by the equation

$$dE_r = \hat{\rho}_s(\theta_i) dE_i \quad (22)$$

where dE_i is the flux density of the incident ray, θ_i is the incident elevation with respect to the local facet, and $\hat{\rho}_s$ is the microscale specular reflection coefficient for the facet.

4.3 Specular Transmission

The method may be used to model microgeometries that include transparent materials. Whenever a ray encounters a

smooth interface between media of different refractive indices, we must calculate the energy transfer through the interface. Neither energy flux density nor radiance is preserved at the interface [4], since solid angles are altered, but the distribution of transmitted rays accounts for this. We also must model any attenuation of the ray as it passes through a transparent medium; for a uniform medium, the ray is attenuated by $e^{-\kappa s}$ where s is the path distance and κ is an extinction coefficient determined by the material.

4.4 Directional-Diffuse Reflection

The most complex transfer takes place when a ray strikes a facet that shows directional-diffuse scattering. When a ray hits such a facet, we send out n rays to the hemisphere above the facet and weight them according to $\hat{\rho}_{bd}$; this serves as a discrete approximation of scattering according to the ideal-diffuse and directional-diffuse parts of the BRDF. The total energy transfer is determined by

$$\begin{aligned} dE_r(\theta_r, \phi_r) &= dI_r d\omega_r \cos \theta_r \\ &= dE_i \hat{\rho}_{bd}(\theta_i, \phi_i, \theta_r, \phi_r) d\omega_r \cos \theta_r \end{aligned} \quad (23)$$

where dE_r is the reflected energy flux density in a particular direction, $\hat{\rho}_{bd}$ is the diffuse part (including directional-diffuse) of the microscale BRDF, and $d\omega_r$ is the solid angle of reflection. The angles (θ_r, ϕ_r) give the reflection direction with respect to the local facet. We multiply by $d\omega_r \cos \theta_r$ to convert the radiance given by $\hat{\rho}_{bd}$ to energy flux density for the next scattering event.

In our implementation reflected rays are cast randomly into the hemisphere above the local (microscale) surface; they are distributed uniformly over this hemisphere, so each ray represents a solid angle of

$$d\omega_r = \frac{2\pi}{n} \quad (24)$$

where 2π is the total solid angle of the hemisphere and n is the number of reflected rays shot.

4.5 Integrating Over Milligeometry

We have described the possible microscale events of a single ray striking a point on the surface, but we must integrate over the specimen to obtain the aggregate BRDF. Just as the Monte Carlo integration was extended in Section 3 to accommodate the two dimensions of the incident hemisphere, it can be extended further to integrate over a two-dimensional specimen surface. We can keep the incident energy flux density constant by keeping both the total incident flux and the receiving surface area constant. We do this by shooting a constant number of rays (energy flux) and by distributing them over a constant surface area. The simplest way to do this is to select a fixed region of the surface, as shown in Figure 4, and to distribute the samples uniformly over this region at each incident angle. The direction of each ray is determined by the incident angles (θ_i, ϕ_i) with respect to the mean surface; its origin will be calculated so that the ray will strike the notional plane of the surface, shown in Figure 4 in red, at the chosen sample point.



Figure 4: Target Area

The surface region chosen should be

- large with respect to the lateral geometric features of the surface, to assure a good statistical average of large-scale scattering;
- large with respect to the *vertical thickness* of the surface geometry; and
- a subset of the total surface geometry, since geometry outside the nominal surface region will be important at high incident angles.

When a ray leaves the specimen area, we update the approximate matrix $\widetilde{\mathbf{M}}$ by adding $\mathbf{Y}(U)\mathbf{Y}^T(V)R/\cos\theta_r$. This matrix represents the BRDF ρ_{bd} . We integrate over the portion of the surface that is visible from the reflection direction (ϕ_r, θ_r) , projected onto the mean surface.

4.6 Efficiency Considerations

We can reduce the computation needed to maintain the matrix $\widetilde{\mathbf{M}}$ by holding the incident direction U constant for several reflected directions V , updating the matrix only once for each distinct U . This happens automatically when several randomly distributed rays are spawned at each intersection, as in directional-diffuse scattering. In addition, we choose several target points on the surface for each U , further amortizing the cost of updating the matrix. Updating the matrix then becomes a triple sum

$$\widetilde{\mathbf{M}} = \frac{1}{N_i N_p N_b} \sum_{n=1}^{N_i} \mathbf{Y}(U_n) \left\{ \sum_{m=1}^{N_p} \sum_{l=1}^{N_b} \mathbf{Y}^T(V_{ml}) \frac{R_{nml}}{\cos\theta_r} \right\} \quad (25)$$

where R_{nml} is the attenuation of a ray from incident direction U_n reflected in direction V_l from target point P_m on the surface. N_i is the number of incident directions used, N_p is the number of sample positions across the specimen for each incident direction, and N_b is the number of exit rays resulting from a single incident ray. This approach reduces the number of evaluations of the spherical harmonics basis functions; for $N_i N_p N_b$ samples to update the matrix, $\mathbf{Y}(U)$ is evaluated only N_i times, while $\mathbf{Y}^T(V)$ is evaluated $N_i N_p N_b$ times. The greatest savings, however, comes in matrix adds; we need only perform N_i matrix additions; the other updates simply add vectors and require far less computation.

4.7 Convergence Measure

Since the exact matrix \mathbf{M} is symmetric, we can use the asymmetry of our estimate as a measure of convergence in approximating the true BRDF. We calculate the error Q as

$$Q = \left\| \widetilde{\mathbf{M}} - \widetilde{\mathbf{M}}^T \right\| \quad (26)$$

where

$$\|\mathbf{A}\| = \frac{1}{N^2} \sum_{j=1}^J \sum_{k=1}^J |\mathbf{A}_{jk}| \quad (27)$$

where J is the size of the matrix \mathbf{A} . This is perhaps not as informative as a direct estimate of the variance of each coefficient, but is much cheaper to compute and tends to decline as $\frac{1}{\sqrt{N}}$, which suggests that it is directly proportional to the variance.

5 Results

We now show several applications of the technique. We obtain BRDF's for surfaces textured at milliscale. At the microscale, the BRDF can be ideal specular, ideal diffuse, or an analytical BRDF that includes wave optics effects. The technique can also be used recursively by using the results of one simulation as the microscale BRDF in another simulation.

All images shown in this section were generated by Monte Carlo ray tracing; the grainy texture of the images is caused by the Monte Carlo integration used to compute global illumination. Other global illumination and rendering techniques might have been used, such as that of Sillion et al. [13].

We first consider a flat Gaussian-rough surface for which, at the microscale, the surface is an ideal specular reflector. We can compare the results of the new method with the results of an existing analytical model for such a surface [5], thus giving some verification of the new technique.

5.1 Initial Verification: An Isotropic Surface

Wave optics effects were not included, except for the Fresnel coefficient for each microfacet. Reflection is governed by geometric optics; shadowing and masking effects of the surface are included because of the occlusion calculations in the ray tracer.

Gaussian height fields were generated by FFT filtering of white noise, and the resulting points were connected by triangles, each of which was modeled as a mirror. To integrate over a specimen large compared to the roughness height, an area of 8×8 millimeters was used. To assure adequate representation of the surface, a total of 524,288 polygons was used. The model was created in four sections of 131,072 polygons, each generated with a different random number seed, to represent a square patch of surface 4mm wide. The roughness length parameters of the surface were $\sigma = 10\mu\text{m}$ vertically and $\tau = 65\mu\text{m}$ horizontally (Figure 5). The specimen patch actually used was 3.13mm wide in the center of the geometric model; this assured that all incident rays would intersect the "sides" of the patch at least 2σ away from the notional plane.

Incident ray angles were restricted to 88° to keep the effective roughness greater than 460 nm, the shortest wavelength employed. This keeps behavior in the regime where geometric optics is valid; were the wavelength to approach the size




Figure 5: Gaussian Surface

of surface features, wave-related effects would begin to affect the scattering. Results are plotted as solid lines in Figure 6 for incident angles $\theta_i = 0^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ$. Dashed lines show results from the model of He [5], which assumes a Gaussian rough surface and allows for wave optics effects. The He model is shown in the limit of large surface roughness, $\sigma \gg \lambda$, where wave optics effects should be negligible. The simulation agrees quite well with the analytical model for reflection angles less than about 80° ; the divergence at greater angles is disturbing, but not very significant in terms of energy values. Recall that the BRDF ρ_{bd} gives a radiance value dI_r ; the energy dE_r scattered in any reflected direction (θ_r, ϕ_r) is proportional to $dI_r \cos \theta_r$, reducing the effect of the error at high angles of reflection. We believe that the error results because we approximate $\rho_{bd} \cos \theta_r$. If we assume that error in approximating this function is roughly constant over the hemisphere, dividing by $\cos \theta_r$ to recover ρ_{bd} will magnify the error near the horizon (i.e. as $\theta_r \rightarrow \frac{\pi}{2}$).

5.2 Simple Anisotropy

We can use the method to create an anisotropic milliscale BRDF by using an isotropic analytical microscale BRDF model; we rely on He's analytical model for microgeometric effects, and use the new technique to model larger-scale anisotropy. Figure 7 shows, at the top, a model of parallel cylinders of slightly rough aluminum. In the left side of the figure, the cylinders are oriented with axes perpendicular to the screen; in the right side the axes are parallel to the screen. The bottom half of the figure shows a similar scene, but with two flat plates replacing the arrays of cylinders. Both plates use a BRDF generated from parallel cylinders like those in the top half of the figure. In the left half of the figure, the axis of anisotropy was oriented perpendicular to the screen; in the right half, it is oriented left-right.

The scattering patterns are similar; when viewed from a distance, the images look the same. The microscale BRDF is important for generating the upper images; the milliscale BRDF is used for the lower figures. Note how the surface orientation affects the appearance, revealing the anisotropic behavior of the reflected light. This is further illustrated in




Figure 6: Comparison with Previous Model

Figure 8, where the same object is rendered with two BRDFs for brushed aluminum, one isotropic and one anisotropic.

Figures 9 and 10 show an aluminum automotive wheel and an aluminum teapot created using this anisotropic BRDF. The polishing scratches were oriented as from rotation, about the vertical axis of the teapot and about the hub of the wheel. The energy-consistency of the BRDF, not guaranteed by previous approaches, allows an accurate global illumination solution.




Figure 7: Parallel Cylinder Model of Anisotropic Surface




Figure 8: Isotropic and Anisotropic Aluminum



Figure 10: Anisotropic Aluminum Teapot



Figure 12: Velvet Doughnut




Figure 11: Microscale Geometry for Velvet




Figure 13: Microscale Structure of Cloth Model

5.3 Velvet

A more complex microgeometry is that of velvet: it consists of many roughly parallel specular fibers extending from a fabric base. This was modeled as a forest of narrow cylinders, with the angle of each cylinder perturbed randomly (Figure 11). The target area for incident rays is shown at the top of the fibers. The fibers are shown as ideal diffuse for clarity; in the BRDF simulation, the fibers were transparent ideal-specular plastic. Whenever a ray intersected a fiber, it was either reflected (with probability equal to the Fresnel reflectivity) or transmitted; when it intersected the base plane, it was absorbed. Figure 12 shows an image made using the resulting BRDF.

5.4 Woven Cloth

The method can also be used recursively to model several scales of roughness; this is demonstrated by modelling woven cloth as shown in Figure 1. At the milliscale, the




Figure 14: Cloth Microscale Geometry and Real Cloth



Figure 9: Anisotropic Aluminum Wheel



Figure 15: Nylon Cushion

weave pattern of the cloth was modeled as shown in Figure 13, and an anisotropic BRDF was used to model the scattering from individual fibers in the threads. The scattering from the surface of each thread (microgeometry) was modeled by the same geometry used in Section 5.2, but using a Fresnel reflectance function to simulate black synthetic fibers. Figure 14 has three parts: on the left, the cloth microgeometry is shown with an ideal-diffuse BRDF; in the center, it is shown with the thread BRDF, and on the right is a magnified photograph of actual cloth. Figure 15 shows a cushion upholstered in black nylon, rendered using the BRDF obtained from this process.

6 Conclusion

Three main points are described in this paper: a new representation of the BRDF, a Monte Carlo technique to estimate the coefficients of the representation, and the means of creating a milliscale BRDF from microscale scattering events. These allow the prediction of scattering for essentially arbitrary geometries. BRDFs for complex surfaces can be simulated hierarchically by using the result of one simulation in generating the BRDF for the next larger scale.

The new representation is concise and well-suited for use in rendering and global illumination calculations. The technique of [13] can be easily extended to accommodate the new representation. Its ease of evaluation suits it for other global illumination methods such as stochastic ray tracing [7, 18] as well.

The Monte Carlo integration used here enables us to model the scattering of many surfaces which have hitherto been impossible to model in computer graphics, producing accurate models for anisotropic surfaces and surfaces with transparent elements.

7 Acknowledgements

The authors would like to thank the Hewlett-Packard Company for donating the Apollo and HP workstations used in this work. Stephen Westin was supported by a fellowship funded by Ford Motor Company. Thanks to Xiao-Dong He for valuable discussions of the physical process of scattering, and to Don Greenberg and Roy Hall for careful reading and criticism of the manuscript. Special thanks to Julie and Kurk Dorsey and to Harold Zatz for their help in assembling the manuscript. Suzanne Smits assisted ably and patiently with the illustrations. Thanks also to Ben Trumbore, Jim Ferwerda, and Hurf Sheldon for maintaining an excellent software and hardware environment. Geometric models were graciously provided by Sabine Coquillart of INRIA for the cushion in Figures 1 and 15, and by Ford Motor Company Design Staff for the wheel in Figure 9.

References

- [1] Cabral, B., Max, N., and Springmeyer, R., Bidirectional reflection functions from surface bump maps. In *Proceedings of SIGGRAPH '87* (July 27-31, 1987, Anaheim, California), *Computer Graphics* 21, 4 (July 1990), 273-281.
- [2] Cook, R. L., and Torrance, K. E. A reflectance model for computer graphics, *ACM Transactions on Graphics* 1, 1 (January 1982), 7-24.
- [3] Gouraud, H. Continuous shading for curved surfaces. *IEEE Transactions on Computers* 20, 6 (June 1971), 623-628.
- [4] Hecht, E., and Zajac, A. *Optics*. Addison-Wesley, 1974.
- [5] He, X. D., Torrance, K. E., Sillion, F. X., and Greenberg, D. P. A comprehensive physical model for light reflection. In *Proceedings of SIGGRAPH '91* (July 28-August 2, 1991, Las Vegas, Nevada), *Computer Graphics* 25, 4 (July 1991), 175-186.
- [6] Goral, C. M., Torrance, K. E., Greenberg, D. P., and Battaile, B. Modeling the interaction of light between diffuse surfaces. In *Proceedings of SIGGRAPH '84* (July 23-27, 1984, Minneapolis, Minnesota), *Computer Graphics* 18, 3 (July 1984), 213-222.
- [7] Kajiya, J. The rendering equation. In *Proceedings of SIGGRAPH '86* (August 18-22, 1986, Dallas, Texas), *Computer Graphics* 20, 4 (August 1986), 143-150.
- [8] Kajiya, J. Anisotropic reflectance models. In *Proceedings of SIGGRAPH '85* (July 22-26, 1985, San Francisco, California), *Computer Graphics* 19, 4 (July 1985), 15-21.
- [9] Kajiya, J., and Von Herzen, B. Ray tracing volume densities. In *Proceedings of SIGGRAPH '84* (July 23-27, 1984, Minneapolis, Minnesota), *Computer Graphics* 18, 3 (July 1984), 165-174.
- [10] Kalos, M. H., and Whitlock, P. A. *Monte Carlo Methods*. John Wiley & Sons, 1986.
- [11] Poulin, P., and Fournier, A. A model for anisotropic reflection. In *Proceedings of SIGGRAPH '90* (August 6-10, 1990, Dallas, Texas), *Computer Graphics* 24, 4 (August 1990), 273-282.
- [12] Phong, B-T. Illumination for computer generated pictures., *Communications of the ACM* 18, 6 (June 1975), 311-317.
- [13] Sillion, F. X., Arvo, J., Westin, S. H., and Greenberg, D. P. A global illumination solution for general reflectance distributions. In *Proceedings of SIGGRAPH '91* (July 28-August 2, 1991, Las Vegas, Nevada), *Computer Graphics* 25, 4 (July 1991), 187-196.
- [14] Siegel, R., and Howell, J. R. *Thermal Radiation Heat Transfer*. Hemisphere Publishing, New York, 1981.
- [15] Torrance, K. E. and Sparrow, E. M. Theory for off-specular reflection from roughened surfaces. In *Journal of the Optical Society of America* 57, 9 (September 1967) 1105-1114.
- [16] Torrance, K. E. and Sparrow, E. M. Off-specular peaks in the directional distribution of reflected thermal radiation. In *Journal of Heat Transfer* (May 1966) 223-230.
- [17] Wallace, P. R. *Mathematical Analysis of Physical Problems*. Dover Publications, Mineola, N. Y. 1984.
- [18] Ward, G. J., Rubinstein, F. M. and Clear, R. D. A ray tracing solution for diffuse interreflection. In *Proceedings of SIGGRAPH '88* (August 1-5, 1988, Atlanta, Georgia), *Computer Graphics* 24, 4 (August 1988), 85-92.
- [19] Whitted, T. An improved illumination model for shaded display. In *Communications of the ACM* 23, 6 (June 1980) 343-349.