

The Halfway Vector Disk for BRDF Modeling

DAVE EDWARDS, SOLOMON BOULOS, JARED JOHNSON and PETER SHIRLEY

University of Utah

MICHAEL ASHIKHMIN

State University of New York, Stony Brook

MICHAEL STARK

University of California, Irvine

and

CHRIS WYMAN

University of Iowa

We present a mathematical framework for enforcing energy conservation in a BRDF by specifying halfway vector distributions in simple two-dimensional domains. Energy-conserving BRDFs can produce plausible rendered images with accurate reflectance behavior, especially near grazing angles. Using our framework, we create an empirical BRDF that allows easy specification of diffuse, specular, and retroreflective materials. We also present a second BRDF model that is useful for data fitting; although it does not preserve energy, it uses the same halfway vector domain as the first model. We show that this data-fitting BRDF can be used to match measured data extremely well using only a small set of parameters. We believe that this is an improvement over table-based lookups and factored versions of BRDF data.

Categories and Subject Descriptors: I.3.7 [Computing Methodologies]: Computer Graphics

General Terms: BRDF, Reflection models

Additional Key Words and Phrases: BRDF, global illumination, Monte Carlo

1. INTRODUCTION

In the absence of significant translucency and florescence, the directional and spectral properties of surface reflection can be described by the *bidirectional reflectance distribution function* (BRDF). An ideal physically-based BRDF is both reciprocal and energy-conserving. Reciprocity is expressed formally using the equation (see Figure 1 for notation):

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \rho(\hat{\omega}_o, \hat{\omega}_i), \quad (1)$$

Author's address: Dave Edwards, University of Utah School of Computing, 50 S. Central Campus Dr., Salt Lake City, Utah, 84112.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0730-0301/20YY/0100-0001 \$5.00




Fig. 1. Notation used in this paper.

and energy conservation is expressed as follows:

$$\forall \hat{\omega}_o, \int_{\Omega_+} \rho(\hat{\omega}_i, \hat{\omega}_o) (\hat{\omega}_i \cdot \hat{N}) d\Omega(\hat{\omega}_i) \leq 1, \quad (2)$$

where ρ is the BRDF. Throughout this paper, we will use hats to denote unit-length vectors, and a dot to indicate the scalar product.

Reciprocity and energy conservation are discussed in detail by Lewis [1993]. In his terminology, BRDFs that are both reciprocal and energy-conserving are called *physically plausible*. It is difficult to create a mathematical expression for a BRDF that satisfies reciprocity, energy conservation, and represents a general-purpose, realistic material. However, many BRDF models produce realistic images; such models are *visually plausible*, even though they violate reciprocity or energy conservation, or both.

Various BRDF models have been developed from first principles using height correlation methods (e.g., [He et al. 1991; Stam 1995]), and microfacet models (e.g., [Cook and Torrance 1982; Oren and Nayar 1994; Ashikhmin et al. 2000]). Other models are phenomenological; high level features such as highlights are described by empirical expressions (e.g., [Phong 1975; Ward 1992; Schlick 1994; Ashikhmin and Shirley 2001]). Some BRDF models use explicit microfacet geometry and simulation to generate BRDF data [Westin et al. 1992].

Another approach is to directly measure the BRDF from samples of real materials using a device such as a gonioreflectometer (e.g., [Marschner et al. 1999; Matusik et al. 2004]). One problem with measured BRDF data is that it has significant statistical and systematic error [Greenberg et al. 1997]. Therefore, it is common to fit measured data to a BRDF model by projecting the data onto one or more smooth curves (e.g., [Lafortune




Fig. 2. Reflective spheres in a constant environment using different BRDFs. Top row: Ward, Ashikhmin-Shirley, Blinn-Phong. Bottom row: Our data-fitting BRDF, Cook-Torrance, our empirical BRDF.

et al. 1997]). Empirical models with intuitive parameters can also be used to fit measured data (e.g., [Ward 1992]). Such models are convenient because they have few parameters, but they may not include subtle features such as retroreflection.

General purpose, physically plausible BRDF models are difficult to formulate, because the constraints imposed by reciprocity and energy conservation are not easily satisfied by a simple mathematical expression. Most empirical BRDF models enforce reciprocity at the cost of energy conservation (e.g., [Lafortune et al. 1997; Ashikhmin and Shirley 2001]). In such models, the lobes of the BRDF may extend below the surface. Therefore, these models predict that some of the light reflected from a surface comes from below the surface, a situation that is physically impossible. Renderers based on such models will attempt to gather light coming from below the surface, which cannot contribute to outgoing light. Visually, this phenomenon appears as “light leakage”: the loss of energy at near-grazing viewing directions causes surfaces to appear darker near silhouette edges.

Other BRDF models produce energy, which also violates energy conservation. For example, any model with a constant diffuse component combined with a Fresnel-varying specular component produces energy, because the specular term goes to unit volume at grazing angles (e.g., [Schlick 1994]). Figure 2 shows spheres rendered using several common BRDFs in a uniform radiance environment. A sphere using an ideal energy-conserving BRDF would not be visible; it would reflect all the light from the environment at each point of the sphere, so every point on the sphere would be the same color as the environment.

For BRDF models in Monte Carlo rendering, it is useful to be able to *importance sample* the BRDF. That is, for a given $\hat{\omega}_o$, it should be practical to generate a directional sample $\hat{\omega}_i$ with density function $Q(\hat{\omega}_i)$ that has the following property:

$$Q(\hat{\omega}_i) \sim \rho(\hat{\omega}_i, \hat{\omega}_o)(\hat{\omega}_i \cdot \hat{N}). \quad (3)$$

Depending on the form of the BRDF model, this is not always possible (see [Ashikhmin

and Shirley 2001]), but in some cases it is (e.g., [Lafortune et al. 1997]). Ease of sampling is usually achieved by making the cosine-weighted BRDF integrable, and making sure that the resulting indefinite integral has certain inversion properties [Pharr and Humphreys 2004].

So far, no general-purpose empirical BRDF model enforces both energy conservation and reciprocity while allowing efficient importance sampling. For this reason, different models are used for different applications, depending on the priorities of the rendering system. Although most BRDF models value reciprocity over energy conservation, both components are important for physically accurate BRDF models. Our main goal in this research is the creation of a framework for BRDF models that enforces energy conservation at the cost of reciprocity. The framework uses simple probability density functions on a bounded, planar domain to describe energy-conserving BRDFs. Using this framework, we have developed a new empirical energy-conserving BRDF model. This model can represent many different reflectance effects using a small number of parameters. We have also developed a similar BRDF model that does not conserve energy but is useful for fitting measured BRDF data. Both new BRDFs can be efficiently sampled.

In the next section, we describe a framework for specifying energy-conserving BRDFs. Section 3 describes a bounded planar domain that can be used with the framework to describe BRDF properties. In Section 4, we describe an empirical energy-conserving BRDF model that we designed using our framework. Section 5 describes another BRDF model that is useful for fitting measured data.

2. A FRAMEWORK FOR ENERGY CONSERVING BRDFs

In this section, we describe a mathematical constraint that will force a BRDF to conserve energy. Initially, we describe this constraint in terms of probability distributions on the set of incident light directions. Ultimately, we show how we can transform the set of incident directions $\hat{\omega}_i$ into other domains while continuing to preserve the constraint.

An energy-conserving BRDF must follow the constraint imposed by Equation 2. For our framework, we will assume the BRDF satisfies an equality instead of an inequality:

$$\forall \hat{\omega}_o, \int_{\Omega_+} \rho(\hat{\omega}_i, \hat{\omega}_o)(\hat{\omega}_i \cdot \hat{N}) d\Omega(\hat{\omega}_i) = 1. \quad (4)$$

This form of the equation allows us to represent a purely reflective surface, which perfectly conserves all light energy. To represent more general materials, we can multiply ρ by a coefficient (i.e., the albedo of the surface) to account for absorption of different wavelengths of light. This ensures that all energy lost is due to light absorption within the surface, a phenomenon that has a physical basis. Therefore, if we can design BRDFs for which the equation holds, we can guarantee that they will conserve energy.

A BRDF has two parameters: the direction toward the incident light $\hat{\omega}_i$, and the direction toward the viewer of outgoing light $\hat{\omega}_o$. Both of these parameters are unit-length direction vectors, and hence lie in a two-dimensional domain; namely, the hemisphere Ω_+ above the surface. For the rest of this discussion, we will assume that $\hat{\omega}_o$ is given (as it is in a Monte Carlo renderer). In this case, the BRDF becomes a function of only $\hat{\omega}_i$. This version of the BRDF is a function defined on the two-dimensional domain Ω_+ .

Since the value of the integral in Equation 4 is 1, we can also view the function:

$$Q(\hat{\omega}_i) = \rho(\hat{\omega}_i, \hat{\omega}_o)(\hat{\omega}_i \cdot \hat{N}), \quad (5)$$

as a probability density function (PDF) over the set of incident directions $\hat{\omega}_i$ on the hemisphere. Rearranging this equation gives us the following form of the BRDF:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{Q(\hat{\omega}_i)}{(\hat{\omega}_i \cdot \hat{N})}. \quad (6)$$

The key to our framework is that as long as $Q(\hat{\omega}_i)$ represents a PDF on the hemisphere Ω_+ , the corresponding BRDF ρ will be energy-conserving. This formulation allows a great deal of freedom, since for a given direction $\hat{\omega}_o$, $Q(\hat{\omega}_i)$ can be any PDF.

However, it can be difficult to design a PDF $Q(\hat{\omega}_i)$ that corresponds to a visually plausible BRDF. For example, phenomena like off-specular reflection cannot be easily formulated in the space of incident directions [Blinn 1977; Ngan et al. 2004]. We now describe a domain in which it is easier to create PDFs that correspond to energy-conserving BRDFs.

Many BRDF models have used the *halfway vector*:

$$\hat{H} = \frac{\hat{\omega}_i + \hat{\omega}_o}{\|\hat{\omega}_i + \hat{\omega}_o\|}. \quad (7)$$

In addition to having a long history of use (e.g., [Torrance and Sparrow 1967; Blinn 1977]), the halfway vector has been used in many recent BRDF models [Rusinkiewicz 1998; Ashikhmin et al. 2000; Ashikhmin and Shirley 2001; Stark et al. 2005]. Ngan et al. [2004] have discussed some underlying reasons for this and have shown empirically that BRDFs based on the halfway vector yield more visually plausible results than BRDFs based on the incident direction $\hat{\omega}_i$. Another advantage of the halfway vector is that in a microfacet-based BRDF model, the vector \hat{H} equals the normal vector of the microfacets responsible for reflection.

Torrance and Sparrow [1967] showed that the PDF $Q(\hat{\omega}_i)$ over incident directions is related to a PDF $q(\hat{H})$ over halfway vectors by the formula:

$$Q(\hat{\omega}_i) = \frac{q(\hat{H})}{4(\hat{\omega}_i \cdot \hat{H})}. \quad (8)$$


Combining Equations 6 and 8 gives us the following form of the BRDF:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{q(\hat{H})}{4(\hat{\omega}_i \cdot \hat{N})(\hat{\omega}_i \cdot \hat{H})}. \quad (9)$$

Once again, as long as $q(\hat{H})$ is a PDF on the set of allowable halfway vectors (i.e., those vectors that for a given $\hat{\omega}_o$ correspond to $\hat{\omega}_i$ in the hemisphere), the corresponding BRDF ρ will conserve energy.

It is easier to describe visually plausible BRDFs using PDFs on the set of halfway vectors than it is using PDFs on the set of incident light directions. In effect, we are specifying a virtual microfacet distribution for each $\hat{\omega}_o$ that corresponds to the reflectance properties of the surface. For many surfaces, the distribution of microfacet normals is more intuitive than the distribution of $\hat{\omega}_i$ for every $\hat{\omega}_o$. An added benefit of the microfacet-based view is that we can adjust the halfway vector PDFs to create the same visual phenomena as the shadowing and masking terms in other microfacet-based BRDF models.

Although halfway vector PDFs are more convenient than PDFs over incident directions, for a given $\hat{\omega}_o$, the set of allowable halfway vectors is a subset of the hemisphere that changes in a complicated way as $\hat{\omega}_o$ approaches a grazing angle. Because of this complexity, it is not easy to define PDFs on the domain of allowable halfway vectors.

Fig. 3. The set of legal halfway vectors \hat{H} .

Ideally, we want to find a domain that allows us to easily define PDFs that correspond to visually plausible, energy-conserving BRDFs. In the rest of this section, we assume that such a domain exists and explain the relation of PDFs to BRDFs on this domain. In the next section, we give a specific example of such a domain.

We define an arbitrary domain D and a bijection $f(\hat{H}) = \vec{h}$ between valid halfway vectors on the hemisphere and points \vec{h} in D . We also define a PDF $p(\vec{h})$ over points in D . We equate differential probabilities between D and Ω_+ using the equation:

$$p(\vec{h})d\mu(\vec{h}) = q(\hat{H})d\Omega(\hat{H}), \quad (10)$$

where $d\mu$ is the differential measure over D . Combining Equations 9 and 10 gives:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{p(\vec{h})d\mu(\vec{h})}{4(\hat{\omega}_i \cdot \hat{N})(\hat{\omega}_i \cdot \hat{H})d\Omega(\hat{H})}. \quad (11)$$

This equation shows that any PDF $p(\vec{h})$ defined on D corresponds to a unique energy-conserving BRDF ρ . This allows us to define an energy-conserving BRDF in terms of a distribution on D . If we choose D intelligently, it should be easier to describe physically plausible BRDFs by reasoning about PDFs in that domain. Since we use halfway vectors as an intermediate domain, a BRDF created using this framework should have the same advantages as BRDFs based on halfway vector distributions.

We sample the BRDF (e.g., for path tracing) by generating a sample point $\vec{h} \in D$ according to $p(\vec{h})$. Then we transform \vec{h} into its equivalent halfway vector \hat{H} using the equation:

$$\hat{H} = f^{-1}(\vec{h}). \quad (12)$$

The inverse of f must exist because f is a bijection. Finally, we transform \hat{H} into its equivalent $\hat{\omega}_i$ using the equation:

$$\hat{\omega}_i = 2(\hat{\omega}_o \cdot \hat{H})\hat{H} - \hat{\omega}_o. \quad (13)$$

We then use $\hat{\omega}_i$ to sample the incident radiance.




Fig. 4. The disk domain viewed in (u, v) space.

3. EXAMPLE TRANSFORMATIONS

In this section, we give an example of a new domain for halfway vector distributions using the framework described in the previous section. The new domain has properties that make it a useful space for halfway vectors. We show two transformations between valid halfway vectors and points in the new domain, and discuss some of the properties of the transformations.

The unnormalized halfway vector \vec{H} can be expressed as:

$$\vec{H} = \hat{\omega}_i + \hat{\omega}_o. \quad (14)$$

Assuming that $\hat{\omega}_o$ is given, we visualize the set of all such \vec{H} by translating the hemisphere Ω_+ , corresponding to all possible $\hat{\omega}_i$, so that the center of its base lies at the tip of $\hat{\omega}_o$ (see Figure 3). Every point on the translated hemisphere corresponds to an unnormalized halfway vector.

We then transform halfway vectors \vec{H} on the hemisphere to points \vec{h} on the base of the hemisphere, shaded in gray in Figure 3. This allows us to represent halfway vector distributions on a unit-radius disk parallel to the surface. Neumann et al. [1999] also showed a projection onto a disk to be useful for BRDF sampling. The disk is a useful domain because it is planar and bounded, and for intelligently chosen transformations, we can easily define PDFs that correspond to visually plausible BRDFs.

Given orthogonal vectors \hat{u} and \hat{v} defining the local orientation of the surface, points on the disk can be uniquely defined by their u - and v -coordinates (see Figure 4). Although points \vec{h} on the disk are three-dimensional, we can think of them as two-dimensional points in (u, v) space by dropping their normal-direction component. This allows us to think of a PDF $p(\vec{h})$ over points on the disk as a two-dimensional PDF in (u, v) space.

Below, we describe two useful transformations between the hemisphere of unnormalized halfway vectors and the disk.

3.1 The scaling projection

The first transformation scales the halfway vector until its tip lies in the base of the hemisphere (see Figure 5). Given an unnormalized halfway vector \vec{H} , the corresponding point




Fig. 5. The scaling projection.

\vec{h} on the base of the hemisphere is:

$$\vec{h} = \frac{(\hat{\omega}_o \cdot \hat{N})}{(\vec{H} \cdot \hat{N})} \vec{H}. \quad (15)$$

We refer to this transformation as the *scaling projection*.

We now solve Equation 10 for the scaling projection. For this domain, $d\mu(\vec{h})$ is simply dA on the disk. Based on the geometry of the vectors, we find:

$$d\Omega(\hat{H}) = \frac{\cos \theta_h d\mu(\vec{h})}{\|\vec{h}\|^2}, \quad (16)$$

where θ_h is the angle between \vec{h} and the surface normal \hat{N} . In the scaling projection, \vec{h} points in the same direction as the unit halfway vector \hat{H} , so:

$$\cos \theta_h = (\hat{H} \cdot \hat{N}) \quad (17)$$

and using Equation 15:

$$\|\vec{h}\|^2 = \frac{(\hat{\omega}_o \cdot \hat{N})^2}{(\hat{H} \cdot \hat{N})^2}. \quad (18)$$

Combining these equations, we find that:

$$\frac{d\mu(\vec{h})}{d\Omega(\hat{H})} = \frac{(\hat{\omega}_o \cdot \hat{N})^2}{(\hat{H} \cdot \hat{N})^3}. \quad (19)$$

Therefore, a PDF $p(\vec{h})$ over the disk defines a corresponding energy-conserving BRDF:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{p(\vec{h})(\hat{\omega}_o \cdot \hat{N})^2}{4(\hat{\omega}_i \cdot \hat{N})(\hat{\omega}_i \cdot \hat{H})(\hat{H} \cdot \hat{N})^3}. \quad (20)$$

In the scaling projection, \vec{h} points in the same direction as \hat{H} . To importance sample the BRDF, we generate a point \vec{h} on the disk according to the PDF $p(\vec{h})$, and normalize to get




Fig. 6. The orthogonal projection.

a unit halfway vector:

$$\hat{H} = \frac{\vec{h}}{\|\vec{h}\|}. \quad (21)$$

We then use Equation 13 to generate an incident direction $\hat{\omega}_i$ to sample the incoming radiance.

The scaling projection has two main benefits. In (u, v) space, the origin corresponds to the surface normal \hat{N} , which is the halfway vector yielding perfect specular reflection. PDFs $p(\vec{h})$ with high values near the origin yield “shiny” BRDFs.

Furthermore, the center of the disk corresponds to a halfway vector in the retroreflective direction. Thus, we can create retroreflective materials by defining PDFs that have high values near the center of the disk.

3.2 The orthogonal projection

The second transformation maps the halfway vector to the disk along the direction of the surface normal \hat{N} (see Figure 6). Given \vec{H} , the corresponding point \vec{h} on the disk is:

$$\vec{h} = \vec{H} - (\hat{\omega}_i \cdot \hat{N})\hat{N}. \quad (22)$$

We refer to this transformation as the *orthogonal projection*.

We now solve Equation 10 for the orthogonal projection. Once again, $d\mu(\vec{h})$ is just dA on the disk. In this case, we have:

$$\frac{d\mu(\vec{h})}{d\Omega(\hat{H})} = \frac{\|\hat{\omega}_i + \hat{\omega}_o\|^2 (\hat{\omega}_i \cdot \hat{N})}{(\hat{\omega}_i \cdot \hat{H})}. \quad (23)$$

The resulting BRDF is:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{p(\vec{h}) \|\hat{\omega}_i + \hat{\omega}_o\|^2}{4(\hat{\omega}_i \cdot \hat{H})^2}. \quad (24)$$

To importance sample the orthogonal projection, we generate a sample point \vec{h} on the disk according to $p(\vec{h})$. From \vec{h} we compute \vec{H} :

$$\vec{H} = \vec{h} + (\hat{\omega}_i \cdot \hat{N})\hat{N}. \quad (25)$$




Fig. 7. Cross sections of BRDF lobes using different projections but the same underlying PDF $p(\vec{h})$. The left image shows lobes using the scaling projection, the right image using the orthogonal projection. In each image, lobes are shown for outgoing angles at 0, 20, 40, and 60 degrees. Lobes from the scaling projection are wide and exhibit primarily off-specular reflection. Lobes from the orthogonal projection are narrow and centered on the direction of perfect reflection. Lobes in both images become much larger as the outgoing angle approaches grazing.

We then normalize \vec{H} to get the unit halfway vector \hat{H} and proceed as when importance sampling the scaling projection.

As in the scaling projection, the origin in (u, v) space corresponds to the normal direction. In general, we found the orthogonal projection more accurate for fitting measured BRDF data than the scaling projection.

3.3 Differences between the projections

Although both projections map halfway vectors into the same domain, the corresponding BRDFs differ. Using the same PDF $p(\vec{h})$ on the disk, the BRDF produced by the scaling projection will have wider, less intense lobes than the BRDF produced by the orthogonal projection. In addition, the lobes in the scaling projection BRDF will exhibit an off-specular peak that has been shown to be beneficial in other models [Torrance and Sparrow 1967]. See Figure 7 to see shape differences between BRDF lobes from the different projections.

4. A NEW EMPIRICAL BRDF

To describe the distribution of projected halfway vectors on the disk, we define a PDF on the disk in (u, v) coordinates. The PDF should allow distribution of points on the disk in a controlled way. Because the disk is finite, the PDF cannot have infinite support. Therefore, we approximate a two-dimensional Gaussian distribution using the following “lump” function:

$$p(\vec{h}) = \frac{n+1}{\pi R^2} \left[1 - \frac{\|\vec{h} - \vec{c}\|^2}{R^2} \right]^n. \quad (26)$$

The equation describes a unit-volume lump centered at \vec{c} with radius R . The exponent n describes the tightness of the lump; if $n = 0$ the function is a uniform distribution, as n approaches infinity, the function approaches a delta function. Figure 8 shows variations in the lump function for different values of R and n .




Fig. 8. Cross sections of the lump PDF. Curves on the left have a radius $R = 1$, and curves on the right have $R = 0.5$. In each graph, the black (lowest) curve has exponent $n = 1$, the red (middle) curve has $n = 5$, and the blue (highest) curve has $n = 10$. Both graphs use the same scale. In general, the function becomes higher and narrower as R decreases or n increases.

The lump function is a convenient PDF for several reasons. By changing \vec{c} , we center the lump at that point on the disk. For a given \vec{c} , we set R to ensure that the entire lump lies within the disk. Changing n controls the variance of the lump's distribution. Since the lump has unit volume, any distribution defined as a sum of such lumps will be a PDF as long as the sum of coefficients is equal to 1.

We make the lump anisotropic by allowing it to have different radii in the \vec{u} and \vec{v} directions. This elliptical lump function is:

$$p(\vec{h}) = \frac{n+1}{\pi R_u R_v} \left[1 - \frac{(\vec{h}_u - \vec{c}_u)^2}{R_u^2} - \frac{(\vec{h}_v - \vec{c}_v)^2}{R_v^2} \right]^n, \quad (27)$$

where R_u and R_v are the radii of the elliptical lump in the \vec{u} and \vec{v} directions, respectively. When $R_u = R_v$, this form of $p(\vec{h})$ is the same as the form in Equation 26.

Note that R_u and R_v need to be set such that the lump stays completely on the disk. If this is not done, the PDF will not have unit volume and the resulting BRDF will not conserve energy.

For rendering purposes, it is useful to generate samples on the disk weighted according to the lump $p(\vec{h})$. This is easiest in polar coordinates; given uniform random numbers (ξ_1, ξ_2) , we find the polar point (r, θ) to be:

$$(r, \theta) = \left(R \sqrt{1 - \sqrt[n+1]{\xi_1}}, 2\pi \xi_2 \right). \quad (28)$$

The resulting sample point \vec{h} in (u, v) space is:

$$\vec{h} = \vec{c} + (r \cos \theta, r \sin \theta) \quad (29)$$

In practice, the easiest way to sample an anisotropic lump is to generate a sample on an isotropic lump with radius $R = 1$ and then scale the resulting sample by R_u in the \hat{u} direction and R_v in the \hat{v} direction to obtain an anisotropic sample.

We construct our empirical BRDF from a Lambertian diffuse component and an energy-conserving specular component of the form in Equation 11. We use the scaling projection

to transform halfway vectors into points on the (u, v) disk because it allows easy representation of retroreflection and off-specular reflection.

We use Schlick's approximation of the Fresnel term [1994] to blend between the diffuse and specular components:

$$F(\theta) = F_0 + (1 - F_0)(1 - \cos \theta)^5, \quad (30)$$

where F_0 is the reflectance at normal incidence. For dielectrics with refractive index η :

$$F_0 = \left(\frac{\eta - 1}{\eta + 1} \right)^2, \quad (31)$$

which ranges from about 0.02 to 0.07 for common dielectrics. Equation 31 works well in practice even for non-dielectrics, such as metals [Cook and Torrance 1982].

For the specular component of our BRDF, we use a weighted sum of two unit-volume lumps, one centered on the surface normal direction and one centered on the retroreflective direction. This allows us to easily control the reflective properties of the surface. Our empirical BRDF uses the scaling projection to easily specify retroreflection and off-specular reflection. Therefore, the corresponding PDF on the disk is:

$$p(\vec{h}) = (1 - k) \frac{n_s + 1}{\pi R_s^2} \left[1 - \frac{\|\vec{h}\|^2}{R_s^2} \right]^{n_s} + k \frac{n_r + 1}{\pi R_r^2} \left[1 - \frac{\|\vec{h} - \omega_{uv}\|^2}{R_r^2} \right]^{n_r}, \quad (32)$$

where $k \in [0, 1]$ is a weighting coefficient that blends between specular reflection and retroreflection, R_s and n_s are the radius and exponent of the specular lump and R_r and n_r are the radius and exponent for the retroreflective lump. In order to constrain the lumps to stay on the disk, R_r must be no greater than 1, and R_s must be no greater than $(1 - \sin \theta_o)$, where θ_o is the angle between $\hat{\omega}_o$ and \hat{N} .

Given $p(\vec{h})$, our empirical BRDF is:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = (1 - F(\theta_o)) \frac{C_d}{\pi} + F(\theta_o) \frac{p(\vec{h})(\hat{\omega}_o \cdot \hat{N})^2}{4(\hat{\omega}_i \cdot \hat{N})(\hat{\omega}_i \cdot \hat{H})(\hat{H} \cdot \hat{N})^3}. \quad (33)$$

The BRDF has seven parameters. C_d is the diffuse color, F_0 is the specular color at normal incidence, and k , R_s , n_s , R_r , and n_r are as explained above. An anisotropic BRDF requires two additional radii parameters for the lumps: one for the specular lump and one for the retroreflective lump. In practice, many of the parameters can be ignored for a given material; for example, if a material does not have a retroreflective component, k can be set to 0 and R_r and n_r are not used. An isotropic material without retroreflection can be modeled using 4 parameters: the colors C_d and F_0 , and the scalars R_s and n_s .

All BRDF parameters can vary with respect to $\hat{\omega}_o$ because even though $p(\vec{h})$ changes, it remains a PDF on the disk. In practice, we assume that C_d , F_0 , and k are constant over all $\hat{\omega}_o$. As $\hat{\omega}_o$ changes, we set R_s and R_r to their maximum allowable values explained above and change n_s and n_r to generate the material properties we desire. In general, if n_s and n_r increase as $\hat{\omega}_o$ approaches grazing, the material will appear shinier; if they decrease, the specular- and retro-reflections will appear blurrier. Changing these parameters with respect to $\hat{\omega}_o$ allows the model to generate some of the same visual effects as the shadowing and masking terms in other halfway vector-based BRDF models.

Our empirical BRDF model allows a great deal of flexibility. It supports diffuse, specular, and retroreflective materials with a small number of parameters. Because of its formu-




Fig. 9. Various anisotropic spheres rendered with our empirical BRDF using an HDR environment map for illumination.




Fig. 10. Left: A road at night rendered using diffuse reflectors. Right: Using retroreflectors modeled by our empirical BRDF.

lation, it is guaranteed to conserve energy: Figure 2 shows a sphere in a constant environment rendered with our BRDF. As expected, the edges of the sphere are invisible because each point on the sphere reflects all of the incident light.

We have rendered several scenes using our BRDF and found that it produces visually plausible results (see Figures 9, 10, and 11). We can also efficiently importance sample our BRDF.

5. DATA FITTING

Our empirical BRDF is based on physical principles, so it should be able to mimic the reflectance properties of real materials. To test this ability, we obtained data from the MIT/MERL BRDF Database [Matusik et al. 2003] for three materials: metallic blue paint, nickel, and plastic. This dataset contains some of the best available measured BRDF data and has been used in several other recent papers. All materials in the dataset are isotropic.




Fig. 11. Top left: Using Lawrence et al.'s factored BRDF at 1225 samples per pixel. Top right: Reference image for the Princeton scene rendered using measured BRDF data at 99225 samples per pixel. Bottom left: Using our data-fitting BRDF with two simple analytic basis functions at 1225 samples per pixel. Bottom right: Using our empirical BRDF with one analytic lobe at 1225 samples per pixel. Insets show a difference between the reference image and the rendered image, where darker portions of the difference image indicate higher disparity. All difference images in the figure are normalized using the same scale factor for improved visibility.

We further assumed that the retroreflective component of each material was negligible, and we attempted to adjust the parameters of our empirical BRDF to fit the measured data.

We found that our empirical model could not fit the data well enough to produce a visually identical rendering. Despite this fact, the output produced by our BRDF looks visually plausible. Figure 11 shows the closest match between the measured data and our BRDF. The figure is based on a similar image from Lawrence et al. [2004]; the only difference is that we reused measured-data materials for the table and the vase (nickel and plastic, respectively) instead of a Cook-Torrance BRDF. This change allowed us to see more clearly how our BRDF fit measured data, especially on the flat tabletop where the




Fig. 12. Comparison of sampling performance. Left: The scene rendered using Lawrence’s factored BRDF at 400 samples per pixel. Right: Using our data-fitting BRDF at 400 samples per pixel. Insets show a closeup of the highlighted region for both models.

viewing vectors are near grazing angles.

After examining the data, we found that it was not consistent with an energy-conserving BRDF. This inconsistency is probably due to measurement noise, especially for angles near grazing where measurement is difficult; this behavior has been noted before with this data set [Lawrence et al. 2004]. It should also be noted that any direct measurement of real reflectance will include effects such as subsurface scattering, which in general cannot be represented by a BRDF model. For example, the material properties of the metallic blue paint are very different from those of plastic and nickel. We assume that the material is a spray paint made from metal flakes suspended in a substrate. Snyder [1998] has noted that when such materials are measured in laboratories, the measurements appear to violate the reciprocity constraint for BRDFs due to subsurface effects.

Since we were unable to fit our empirical model closely to the data, we built a software tool to examine the distributions of halfway vectors on the disk using either the scaling projection or the orthogonal projection. The tool allows us to examine how the halfway vector distribution changes as $\hat{\omega}_o$ changes. To fit our BRDF to the measured data, we tried both nonlinear optimization and hand fitting using our software tool mentioned above. In general, we could fit the data much better by hand, and after practice with the fitting tool, we were able to perform by-hand fits almost as fast as the optimization software.

We found the distribution for all three materials shows similar behavior: it is approximately constant as $\hat{\omega}_o$ goes from normal incidence to the point where θ_o is approximately $\frac{\pi}{4}$. From that point until $\hat{\omega}_o$ approaches grazing angles, the distribution tightens around the normal direction. See Figure 13 for screenshots of our fitting software showing the distribution of halfway vectors for a plastic material for different values of $\hat{\omega}_o$.

Based on the behavior we saw in the halfway vector distributions, we developed a new BRDF model to better fit the measured data. None of the measured materials show perfect reflectance at grazing angles, so for our data-fitting BRDF, we changed the Fresnel term to




Fig. 13. Screenshots of our BRDF-fitting software showing the distribution of halfway vectors of a plastic material on the disk for different values of $\hat{\omega}_o$. Data points from the measured data are shown in red and points from our data-fitting model are in blue. From top to bottom, the angle between the surface normal and $\hat{\omega}_o$ is 0, 15, and 30 degrees for the left-hand images, and 45, 60, and 75 degrees for the right-hand images. Note that the disk shifts from image to image because it is centered on $\hat{\omega}_o$ in (u, v) space.

allow blending between a specular color and another color (not necessarily white):

$$F_{fit}(\theta_o) = F_0 + (F_{\frac{\pi}{2}} - F_0)(1 - \cos \theta_o)^5 \quad (34)$$

We also found that using the orthogonal projection made it easier to fit the measured data. Instead of fitting a specular lump and a retroreflective lump, we used a PDF based on two specular lumps with different exponents. Summing two such basis functions gave us additional flexibility in fitting and is consistent with the approach found in other data-fitting methods [Lafortune et al. 1997; Lawrence et al. 2004]. Ngan et al. [2004] has shown that using a second such term in a BRDF model can reduce fitting error by more than 20%.

To fit the measured data, we used the following PDF on the disk:

$$p(\vec{h}) = (1-k) \frac{n_1+1}{\pi R_1^2} \left[1 - \frac{\|\vec{h}\|^2}{R_1^2}\right]^{n_1} + k \frac{n_2+1}{\pi R_2^2} \left[1 - \frac{\|\vec{h}\|^2}{R_2^2}\right]^{n_2}. \quad (35)$$

As in the empirical model, $k \in [0, 1]$ is a weighting coefficient; it blends between the two specular lumps. R_1 and n_1 are the radius and exponent of the first specular lump and R_2 and n_2 are the radius and exponent for the second specular lump. We no longer constrain the values of R_1 and R_2 to keep the lumps on the disk. This leads to energy loss in our BRDF, but similar energy loss is apparent in the data.

Given $p(\vec{h})$, the BRDF we used to fit measured data is:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = (1 - F_{fit}(\theta_o)) \frac{C_d}{\pi} + F_{fit}(\theta_o) \frac{p(\vec{h}) \|\hat{\omega}_i + \hat{\omega}_o\|^2}{4(\hat{\omega}_i \cdot \hat{H})^2}. \quad (36)$$

As with our empirical model, the parameters can vary with respect to $\hat{\omega}_o$. We left most parameters constant; however, we achieved the best results by increasing the larger of n_1 and n_2 as $\hat{\omega}_o$ approaches grazing, and by changing k to increase the importance of the corresponding lump. Although this BRDF model does not conserve energy, in practice it still loses less energy than many other empirical models. Figure 2 shows an example of our data-fitting BRDF in a constant environment.

Figure 11 shows the results of fitting this BRDF to the measured data. We compare our results to an image using the measured data directly and an image using a factored BRDF representation. Although our method uses only a few parameters, our image is as good as the table-based factored representation and requires less storage.

Overall, we found that fitting distributions in halfway vector space was simple. Despite their different properties at normal incidence, all three materials showed similar behavior toward grazing angles, which is consistent with optical theory. Furthermore, using a sum-of-basis-functions approach allowed us to better fit the measured data. Theoretically, we could improve our fit by adding more lumps to sample the disk; however, results with only two lumps are excellent.

6. CONCLUSION AND FUTURE WORK

We have presented a framework for transforming halfway vectors into different domains such that the corresponding BRDF will conserve energy. We have shown that a unit-radius disk is one such domain that allows us to easily specify certain BRDF properties such as specular- and retro-reflection. Using our framework, we have created an empirical energy-preserving BRDF that exhibits visually plausible behavior. Because of the energy conservation constraint, our model works especially well near grazing angles, where other models exhibit artefacts due to energy loss. Our empirical BRDF allows the user to specify diffuse, specular, and retroreflective materials with a small number of intuitive parameters.

We also presented a BRDF model useful for data fitting; although it does not preserve energy, it uses the same halfway vector distribution domains as the first model. This data-fitting BRDF can be used to match measured data extremely well using a small number of parameters. Although the space required by our model is a small fraction of that required by measured data or factored representations, we can still produce images that closely match measured BRDFs for a variety of materials.




Fig. 14. Left: Scene rendered using our data-fitting BRDF (also shown in Figure 11). Right: The scene rendered using our data-fitting BRDF with $\hat{\omega}_i$ and $\hat{\omega}_o$ exchanged.

While producing our mathematical framework, we realized that the ability to examine halfway vector distributions in different domains is a powerful tool. In the future, we plan to investigate domains other than the disk. It would be interesting to find a single domain or a small set of domains that allow us to describe halfway vector distributions for all types of materials.

We will also investigate different distribution functions (i.e., other than our lump function) to see if new functions can better approximate more complex materials such as fabrics. Of course, different functions might work better in different domains: for example, if we transform halfway vectors into an infinite two-dimensional domain, we should investigate distribution functions with infinite support.

Our current empirical BRDF does not enforce reciprocity. We have found that although this gives noticeable differences near grazing angles when the incident and outgoing directions are swapped, the final results of our model are still visually plausible (see Figure 14). Nonetheless, it would be useful to develop a framework for enforcing reciprocity as we have done with energy conservation.

Finally, we removed the energy conservation constraint from our BRDF in order to fit measured data. Our data-fitting BRDF exhibits energy loss because our distribution functions stray outside the domain of legal halfway vectors for certain outgoing directions $\hat{\omega}_o$. One way to compensate for this would be to scale the distribution functions to make up for the lost volume over the disk. We hope to find a simple way to approximate this lost volume so that we can enforce energy conservation with our data-fitting BRDF.

REFERENCES

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based brdf generator. 65–74.
- ASHIKHMIN, M. AND SHIRLEY, P. 2001. An anisotropic Phong BRDF model. *Journal of Graphics Tools* 5, 2, 25–32.
- BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *Proceedings of SIGGRAPH*. 192–198.

- COOK, R. L. AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics 1*, 1 (Jan.), 7–24.
- GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P. S., ARVO, J. R., FERWERDA, J. A., PATTANAİK, S., LAFORTUNE, E. P. F., WALTER, B., FOO, S.-C., AND TRUMBORE, B. 1997. A framework for realistic image synthesis. In *Proceedings of SIGGRAPH*. 477–494.
- HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. In *Proceedings of SIGGRAPH*. 175–186.
- LAFORTUNE, E. P. F., FOO, S.-C., NCE, K. E. T., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH*. 117–126.
- LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient brdf importance sampling using a factored representation. *ACM Transactions on Graphics (SIGGRAPH 2004)*.
- LEWIS, R. 1993. Making shaders more physically plausible. In *Eurographics Workshop on Rendering*. 47–62.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based brdf measurement including human skin. In *Eurographics Rendering Workshop*.
- MATUSIK, W., LOPER, M., AND PFISTER, H. 2004. Progressively-refined reflectance functions from natural illumination.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. Efficient isotropic brdf measurement. In *Eurographics Symposium on Rendering*. 241–248.
- NEUMANN, L., NEUMANN, A., AND SZIRMAY-KALOS, L. 1999. Compact metallic reflectance models. *Computer Graphics Forum 18*, 3 (Sept.), 161–172.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2004. Experimental validation of analytical brdf models. *SIGGRAPH Technical Sketch*.
- OREN, M. AND NAYAR, S. K. 1994. Generalization of lambert’s reflectance model. In *Proceedings of SIGGRAPH*. 239–246.
- PHARR, M. AND HUMPHREYS, G. 2004. *Physically Based Rendering*. Morgan Kaufmann.
- PHONG, B.-T. 1975. Illumination for computer generated images. *Communications of the ACM 18*, 6 (June), 311–317.
- RUSINKIEWICZ, S. M. 1998. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop*. 11–22.
- SCHLICK, C. 1994. An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum 13*, 3, 233–246.
- SNYDER, W. C., WAN, Z., AND LI, X. 1998. Thermodynamic constraints on reflectance reciprocity and kirchhoff’s law. *Applied Optics 37*, 16, 3464–3470.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop*. 41–50.
- STARK, M. M., ARVO, J., AND SMITS, B. 2005. Barycentric parameterizations for isotropic BRDFs. *IEEE Transactions on Visualization and Computer Graphics 11*, 2, 126–138.
- TORRANCE, K. E. AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *57*, 9, 1105–1114.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH*. 265–272.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of SIGGRAPH 92*. 255–264.

Received June 2005; revised Month Year; accepted Month Year