



Directional Statistics BRDF Model

Ko Nishino

Department of Computer Science, Drexel University
Philadelphia, PA

kon@drexel.edu

Abstract

We introduce a novel parametric BRDF model that can accurately encode a wide variety of real-world isotropic BRDFs with a small number of parameters. The key observation we make is that a BRDF may be viewed as a statistical distribution on a unit hemisphere. We derive a novel directional statistics distribution, which we refer to as the hemispherical exponential power distribution, and model an isotropic BRDF with a mixture of it. The novel directional statistics BRDF model allows us to derive a canonical probabilistic method for estimating its parameters including the number of components. We show that the model captures the full spectrum of real-world isotropic BRDFs with accuracy comparable to non-parametric models but with a much more compact representation. We also experimentally show that the model achieves better accuracy with less measurements compared with such non-parametric models. We further demonstrate the advantages of the novel BRDF model by showing its use for reflection component separation and for exploring the space of isotropic BRDFs.

1. Revisiting Parametric BRDF Modeling

The appearance of a real-world object is determined by the complex interplay of the illumination, the geometric structure, and the reflectance (material) property of the object surface. If we assume strictly local light interaction, i.e., if we ignore all global transports such as subsurface scattering, we may mathematically model this light interaction with the Bidirectional Reflectance Distribution Function (BRDF) [21]. The BRDF is arguably the most widely used light transport model in computer vision since it offers a pointwise description of light reflection. As such, devising a compact yet accurate representation of real-world BRDFs goes to the heart of a great number of computer vision applications, especially those concerned with the recovery of illumination, material property, or geometry from images. In particular, low-dimensional parametric BRDF models play vital roles since they enable the formulation of such inverse

problems as parameter estimation. Yet, deriving a parametric BRDF model that can accurately describe the drastic variation real-world BRDFs take on remains a challenging problem. In this paper, we specifically focus on isotropic BRDFs.

Various parametric BRDF models have been introduced in the past. These models are either based on purely empirical observation, such as Lambertian [15], Phong [23, 4], and Schlick [27] models; physically-based modeling of the microscopic surface geometry, including Torrance-Sparrow [6, 30] Ward [31], and Oren-Nayar [19] models; or phenomenological modeling using linear/nonlinear bases such as Zernike polynomials [13], spherical harmonics [24], cosine lobes [14], and 2D Gaussians on halfway disks [9]. Unfortunately, each parametric BRDF model is limited to representing only a specific type of reflection, e.g., glossiness around the reflection vector [14, 23, 4] or the halfway vector [6, 30, 31, 9], and cannot express the whole spectrum of BRDFs in a single parametric form; they require a (linear) combination of separate analytical models leading to complex expressions whose parameter estimation becomes challenging [9]. Although phenomenological models aim to represent all BRDFs with a common set of bases, the use of generic bases inevitably lead to high-dimensional representations for real-world BRDF that can have various frequencies and shapes.

Recently, owing to the development of novel gonioreflectometers (e.g., [8, 17, 11]), extensive measurements of real-world BRDFs have been collected. These data sets have inspired the use of various non-parametric BRDF models [2, 8, 16, 25, 28] which essentially provide tabulated views of the measured BRDF data accessible with two to four-dimensional indices that encode the combination of incident and exitant directions. Since non-parametric models are essentially raw measurements, they undoubtedly have strong advantages in photorealistic appearance synthesis. Yet, when using such non-parametric models for solving inverse problems in computer vision, we are cursed by the high-dimensionality of BRDF data. Although compression techniques including linear/nonlinear dimensionality reduction can help in achieving a lower-dimensional non-parametric model, they hardly result in a few dimensions that

can make solving inverse problems tractable, nor they provide physically meaningful descriptions of the BRDF. Furthermore, the accuracy of non-parametric models essentially depend on the sampling of the data and thus necessitates very dense measurements to achieve certain accuracy.

We tackle the challenging problem of deriving a low-dimensional parametric BRDF model that can achieve accuracy comparable to non-parametric BRDF models. We view the BRDF as a directional statistics distribution, a probability density function that takes in an incident light ray direction and returns a distribution of reflected light ray directions¹. To this end, we derive a novel hemispherical directional statistics distribution and model real-world isotropic BRDF with a mixture of it. The novel directional statistics distribution model can model the whole spectrum of BRDFs ranging from purely Lambertian to perfect mirror reflection in the exact same functional form. This allows us to capture various types of real-world isotropic BRDFs in a low-dimensional parametric form. The statistical model also enables the derivation of a canonical probabilistic method for estimating its parameter values and the optimal number of mixture components. Experimental results show that the model achieves accuracy comparable to the state-of-the-art non-parametric model [25] with a much smaller footprint.

We believe the novel directional statistics BRDF model has direct implications for a broad range of applications. We demonstrate this by showing that the model can achieve higher accuracy than non-parametric models when the sampling of measurements is sparse; it can naturally decompose real-world measured BRDFs into physically meaningful and intuitive constituents (reflection components); and it also provides powerful handles for exploring the entire space of isotropic BRDFs.

2. Isotropic BRDF

The BRDF is defined as the ratio of the reflected differential radiance dL_o in a given exitant (view) direction ω_o to the incident irradiance dE_i due to light from direction ω_i ,

$$f_r(\omega_i, \omega_o) = \frac{dL_o(\omega_o)}{dE_i(\omega_i)} = \frac{dL_o(\omega_o)}{L_i(\omega_i)(\omega_i \cdot \mathbf{n})d\omega_i}, \quad (1)$$

where \mathbf{n} is the surface normal at the surface point of interest and $d\omega_i$ is the differential solid angle the light source in direction ω_i subtends [21]. The BRDF is thus a four-dimensional real-valued function $f_r : \Omega \times \Omega \rightarrow \mathbb{R}$, where Ω is the upper hemisphere with its origin at the surface point and its north pole (Z axis) aligned with the surface normal. Real-world BRDFs satisfy the Helmholtz reciprocity $f_r(\omega_i, \omega_o) = f_r(\omega_o, \omega_i)$ and the energy conservation law $\int_{\Omega} f_r(\omega_i, \omega_o)d\omega_i \leq 1$.

¹Note that this is completely different from using non-parametric [2] or analytical distributions to model microfacet orientations (e.g. [30]). We directly model scattered radiance distributions with parametric directional distributions resulting in a compact yet flexible BRDF model.

The two directions ω_i and ω_o can be described in spherical coordinates (θ_i, ϕ_i) and (θ_o, ϕ_o) , respectively, leading to an explicit 4D notation $f_r(\theta_i, \phi_i, \theta_o, \phi_o)$. Here we denote the polar angle with θ and azimuth angle with ϕ : $\theta \in [0, \frac{\pi}{2}]$ and $\phi \in [0, 2\pi)$. Note that in the above definition and throughout the rest of the paper, we abuse the notation for directional vectors: ω_i (or \mathbf{n}) represents a unique directional vector (unit vector) in either spherical coordinates (θ_i, ϕ_i) or Cartesian coordinates $(\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$, whichever appropriate depending on the context².

Although the BRDF is a four-dimensional function, for many real-world materials, its intrinsic dimensionality is less than four. We may exploit this lower-dimensional intrinsic characteristic of real-world BRDF by making a few realistic assumptions and a coordinate change.

First, since we focus on isotropic BRDFs, we may safely assume that the BRDF is invariant to azimuthal rotations of the incident and exitant directions $f_r(\theta_i, \phi_i, \theta_o, \phi_o) = f_r(\theta_i, \phi_i + \varphi, \theta_o, \phi_o + \varphi)$ and to reflection by the incident plane $f_r(\theta_i, 0, \theta_o, \phi_o) = f_r(\theta_i, 0, \theta_o, -\phi_o)$. These two properties lend themselves to a three-dimensional description of the BRDF $f_r(\theta_i, \theta_o, |\phi_i - \phi_o|)$ [28]. Second, we leverage a common reparametrization known as the halfway vector representation [26]. Instead of representing the BRDF with incident and exitant directions, we model them with the halfway direction ω_h between the incident and exitant directions and encode the incident light direction in a hemisphere where the halfway direction becomes the north pole, referred to as the difference direction ω_d . This reparametrization can be achieved with

$$\omega_h = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|} \quad \omega_d = \mathbf{R}_Y(-\theta_h)\mathbf{R}_Z(-\phi_h)\omega_i, \quad (2)$$

where \mathbf{R}_Y and \mathbf{R}_Z denote rotation matrices ($\in \mathbb{SO}(3)$) about the binormal (Y) and surface normal (Z) axes, respectively [26].

With the halfway representation, an isotropic distribution around the halfway vector, i.e., axially symmetric distribution about $\theta_h = 0$, results in an anisotropic distribution around the reflection vector ($\omega_r = 2(\omega_h \cdot \omega_i)\omega_h - \omega_i$) given an incident light direction, whose elongation along the great circle joining the incident and reflection directions increases as the incident direction approaches the grazing angle [20]. This is a particularly useful property since many real-world isotropic BRDFs exhibit asymmetric reflection around the reflection direction that cannot be captured with isotropic distributions around ω_r [23, 14, 4], but can be approximated with an isotropic distribution around their corresponding halfway vectors ω_h . We empirically found that for BRDFs that can be well-approximated with the Torrance-Sparrow reflection model [30], the center-peaked symmetry

²For instance, the dot product in Eq. 1 is defined over the Cartesian coordinates.

(isotropy about $\theta_h = 0$) can be assumed for a wide range of surface roughness: roughly $[0, 0.4]$ in the roughness parameter value, where it generally ranges from 0 to 0.1 for real-world surfaces. This means that we may represent isotropic BRDFs as function(s) of (θ_h, θ_d) by assuming invariance to azimuthal angles ϕ_h or ϕ_d ³. Recently, Romeiro et al. empirically showed that most isotropic BRDFs can indeed be represented with a two-dimensional (non-parametric) function of θ_h and θ_d by simply averaging over ϕ_d [25].

To this end, we will model real-world isotropic BRDFs with a two-dimensional analytic model as we detail next. We seek to estimate an optimal⁴ parametric representation for the 2D slices of the real-world isotropic BRDFs encoded in measured data of the 3D function $f_r(\theta_h, \phi_d, \theta_d)$.

3. Hemispherical Exponential Power Dist.

We will now derive an analytical reflection model for representing real-world isotropic BRDFs. In particular, we will derive a parametric model for $f_r(\theta_h, \phi_d, \theta_d)$, where the azimuthal dependency on the incident-exitant directions are represented with ϕ_d . Note that we may instead choose to model $f_r(\theta_h, \phi_h, \theta_d)$ without any change in the following derivation. Given the results of the analysis in the previous section, we model this three-dimensional distribution that is intrinsically two-dimensional ($f_r(\theta_h, \theta_d)$) with a set of 2D slices of fixed difference directions θ_d (relative incident direction) and $\phi_h = 0$: $f_r(\theta_h, \phi_d, \theta_d = \theta_d^i)$, where $\theta_d^i = \{\frac{m\pi}{2M} | m = 0, \dots, M\}$. Figure 1(a) shows the 1D profiles on the incident plane of $\theta_d = 0$ -slices of several measured BRDFs [16]— $f_r(\theta_h, \phi_d = \{0, \frac{\pi}{2}\}, \theta_d = 0)$. They appear center-peaked and symmetric about $\theta_h = 0$, as predicted by our previous analysis. They also show that different BRDFs can have very different distribution shapes.

The key observation we make is that a 2D slice of a BRDF can be viewed as a statistical distribution of reflected light rays given an incident light ray (θ_d), where the reflected radiance values $f_r(\theta_h, \phi_d, \theta_d)$ represent the probability of incident light being scattered into that specific direction (θ_h, ϕ_d). In other words, we may view it as a probability density function of a directional distribution—a distribution on the surface of a unit hemisphere Ω parameterized with (θ_h, ϕ_d) . Our goal is to derive a suitable directional statistics distribution model that has a small number of parameters while realizing the necessary flexibility to represent a wide-range of real-world isotropic BRDFs.

The analysis in the previous section suggests that we may approximate these BRDF slices with a center-peaked, i.e., mean direction at $\theta_h = 0$, isotropic probability density function of θ_h : $p(\theta_h | \theta_d)$. Conventional directional statistics distributions, such as the von Mises-Fisher distribution [10], a

Gaussian distribution on the unit sphere, lack the flexibility to model a wide variety of directional distributions that real-world BRDFs exhibit. An even bigger problem of such directional statistics distributions is that the domain is the entire sphere, making it unsuitable for modeling a BRDF.

To this end, we derive a novel hemispherical distribution model analogous to the exponential power distribution in Cartesian coordinates. We will refer to this directional distribution as the *Hemispherical Exponential Power Distribution* (hemi-EPD):

$$p(\theta_h | \theta_d, \Theta) = C(\Theta) (\exp[\kappa \cos^\gamma \theta_h] - 1), \quad (3)$$

where $\Theta = \{\kappa, \gamma\}$ are the parameters of the distribution. C is the normalization factor which can be shown to be $C(\Theta) = \frac{\gamma(-\kappa)^{\frac{1}{\gamma}}}{2\pi \left(\Gamma(\frac{1}{\gamma}) - \Gamma(\frac{1}{\gamma}, -\kappa) - \gamma(-\kappa)^{\frac{1}{\gamma}} \right)}$, where Γ is the incomplete gamma function [1].

This hemi-EPD has various advantages for modeling real-world BRDFs. It can naturally encode axially distributed directions with a small number of parameters, while retaining the flexibility to represent a wide variety of distributions on the hemisphere. We will refer to κ as the scale parameter and γ as the shape parameter. In essence κ controls the overall height of the distribution, corresponding to the albedo of the reflected distribution, and γ controls the kurtosis of the distribution, the concentration of the reflected light directions. As extreme cases, the hemi-EPD can model a perfect Lambertian reflection with $\gamma = 0$ and a perfect mirror reflection with $\gamma = \infty$. We may also see that, by taking Maclaurin expansion $p(\theta_h | \theta_d, \Theta) = C(\Theta)(e^{\kappa \cos^\gamma \theta_h} - 1) + O(\theta_h^n)$, the hemi-EPD *subsumes* a Gaussian distribution of θ_h and thus the Torrance-Sparrow model.

4. Mixture of Hemi-EPDs

Real-world object surfaces are rarely made of a single material; at a microscopic level the surface usually consists of multiple layers of materials. Even when subsurface scattering is negligible, as we assume, individual layers contribute to different distributions of scattered exitant radiance. As a result, real-world BRDFs tend to exhibit a complex distribution that cannot be modeled with a single analytical model. To represent such compound radiance distributions, real-world BRDFs are often modeled with a linear combination of multiple parametric BRDF models. For instance, the Lambertian and the Torrance-Sparrow reflection models are often used together to model real-world surfaces [18]. In many cases, however, the glossy appearance itself, or the diffuse appearance itself, is generated from multiple surface materials and thus cannot be captured with a single parametric model. Cook and Torrance [6], for instance, suggest using multiple microfacet distributions for modeling multiple layers of glossy surface material.

³Note that these two azimuthal angles are uniquely interrelated given an incident light direction or exitant direction.

⁴In a least square sense.

The problem of taking linear combinations of different parametric BRDF models is that the resulting model becomes unnecessarily complex when the functional forms are not the same, e.g., Lambertian and Torrance-Sparrow. This becomes a large problem when we try to estimate their parameters in computer vision applications. On the other hand, when the constituent BRDFs models are the same, the resulting BRDF model would lack the expressiveness to represent a wide variety of materials.

We model a real-world isotropic BRDF with a mixture of hemi-EPDs. We do not suffer from the aforementioned dilemma, since the hemi-EPD itself offers the flexibility to model a wide range of radiance distributions in the same functional form. A 2D slice (of a unique θ_d) of an isotropic BRDF is modeled as a linear combination of hemispherical EPDs

$$p(\theta_h|\theta_d, \Theta) = \sum_{k=1}^K \alpha^{(k)} p(\theta_h|\theta_d, \Theta^{(k)}), \quad (4)$$

where $\Theta = [\Theta^{(1)} \dots \Theta^{(K)}]$ and K is the number of constituent hemi-EPDs. Since the hemi-EPDs are axially symmetric directional distributions, the dependency on ϕ_d is implicit (we find least-square optimal fits). We refer to each constituent hemi-EPD as a BRDF lobe. This parametric model can be seen as a mixture of directional distributions and readily provides a statistical interpretation of the BRDF; it is the directional probability density function of the exitant radiances given an incident light ray.

Measured BRDFs are, however, usually not normalized, i.e., the total energy is not one. To model such real-world data, we fit a mixture of unnormalized hemi-EPDs

$$p(\theta_h|\theta_d, \Theta) = \sum_{k=1}^K \frac{1}{C(\Theta^{(k)})} p(\theta_h|\theta_d, \Theta^{(k)}). \quad (5)$$

to them. Note that we may interpret the reciprocal of the normalization factors $\frac{1}{C(\Theta^{(k)})}$ to be the unnormalized mixture coefficients. Whenever necessary, we can scale these unnormalized mixture coefficients to arrive at a valid probability density function, which simply corresponds to scaling the measured BRDF data with $\sum_{k=1}^K C(\Theta^{(k)})$ to normalize its total energy. This in turn means that as long as $\sum_{k=1}^K \frac{1}{C(\Theta^{(k)})} \leq 1$, the energy conservation law is satisfied.

We may derive a canonical algorithm for fitting this mixture of hemi-EPDs, which we refer to as the *directional statistics BRDF model* or in short the *DSBRDF model*, to measured data by following the same principal as that of the Expectation Maximization algorithm [3]. As outlined in Algorithm 1, this amounts to iterating between computing the conditional expectation of the latent variable k (E-step), i.e., estimating the *responsibilities*, and then maximizing the complete joint likelihood (M-step), i.e., maximum likelihood estimation of the parameters of each unnormalized k th

Algorithm 1 EM for Mixtures of Hemi-EPDs

Initialize all $\kappa^{(k)}, \gamma^{(k)}$ $k = 1, \dots, K$

repeat

 EXPECTATION STEP

for $n = 1$ to N **do**

for $k = 1$ to K **do**

$$q(k|\theta_h^{(n)}, \Theta_{\text{old}}) = \frac{\frac{1}{C(\Theta_{\text{old}}^{(k)})} f(\theta_h^{(n)}|\Theta_{\text{old}}^{(k)})}{\sum_{k=1}^K \frac{1}{C(\Theta_{\text{old}}^{(k)})} f(\theta_h^{(n)}|\Theta_{\text{old}}^{(k)})}$$

end for

end for

 MAXIMIZATION STEP

 Initialize $\Theta^{(k)}$ for $k = 1, \dots, K$

for $k = 1$ to K **do**

$$\underset{\Theta^{(k)}}{\operatorname{argmin}} \sum_{n=1}^N \|q(k|\theta_h^{(n)}, \Theta_{\text{old}}) L_i(\theta_h^{(n)}, \phi_d^{(n)}, \theta_d) - \frac{1}{C(\Theta^{(k)})} p(\theta_h^{(n)}|\theta_d, \Theta^{(k)})\|^2$$

end for

until convergence

lobe. The input to the algorithm is a 2D slice of measured BRDF data ($\theta_d = \text{const}$) consisting of N data points, each having different $(\theta_h^{(n)}, \phi_d^{(n)})$. We formulate the maximum likelihood estimation as least square minimization which we solve with the Levenberg-Marquardt algorithm. For the M-step, the parameter values can be initialized by first estimating κ and then using it to estimate γ for each lobe. We use the same strategy for initializing all values in the beginning: we first estimate one set of κ and γ from all the data points and then halve those values as we step through the remaining lobes. We found this initialization to work well for all 100 BRDFs that we fit.

Once we fit the DSBRDF model to the 2D slices of a measured BRDF data ($\{\theta_d^i | i = 0, \dots, M\}$), we can establish correspondences among the K individual lobes across different θ_d^i by simply sorting the lobes based on the scale parameter values $\kappa^{(k)}$, since we may safely assume that individual BRDF lobes will decay or increase coherently—their height ordering will remain in tact as θ_d increases. Note again that we do not simply ignore the dependency of the BRDF slice on ϕ_d ; through the EM algorithm, we estimate an optimal fit of DSBRDF model in a least square sense, i.e., estimate the parameter values that explains the asymmetry with minimum squared error. If necessary, we may also conduct robust estimation by using a robust estimator in the M-step, which may become crucial for rejecting outliers that are often found in real-world measured data.

5. Determining the Number of Lobes

The number of necessary lobes for accurately encoding the directional distribution of a real-world isotropic BRDF varies depending on the material. We can automatically determine the optimal number of lobes by running the EM al-

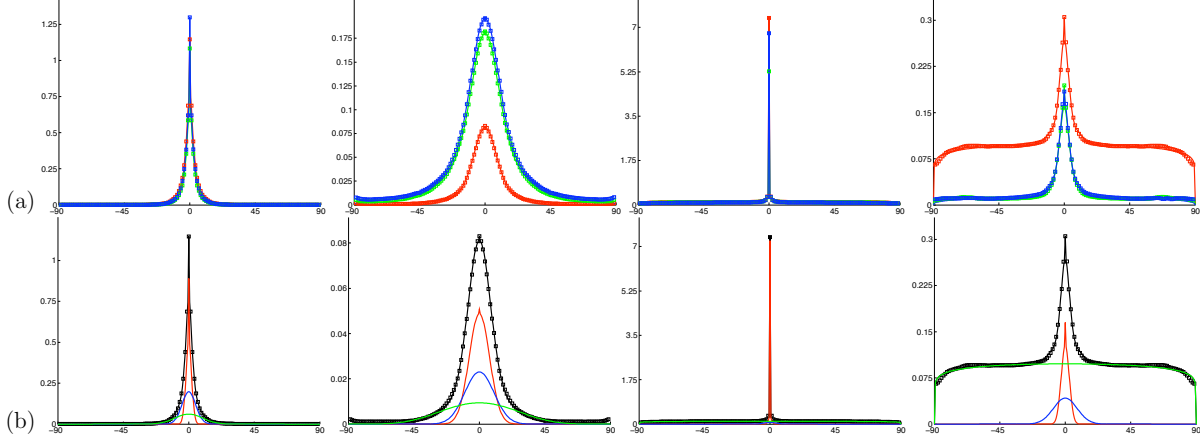


Figure 1: Row (a): The directional statistics BRDF model with 3 lobes (solid line) fit to $\theta_d = 0$ slices of different measured BRDF data [16] (square) shown as the 1D profile on the incident plane. The DSBRDF model accurately fits the measured data despite the significantly varying shapes of the measured distributions (RGB correspond to RGB color channels). Row (b): The lobes (red, green, blue lines) of a 3-lobe DSBRDF model (black solid line) fit to measured data (black dots). Each lobe clearly captures a distinct characteristic reflectance component of the BRDF (zoom in to see the small lobes in the 3rd column BRDF).

gorithm for different numbers of lobes and then testing the quality of fits using a statistical measure.

If we consider the KL-divergence between the statistical distributions of a measured BRDF $f_r(\theta_h, \phi_d, \theta_d)$ and the fit K -lobe directional statistics BRDF distribution $p(\theta_h|\theta_d, \Theta^K)$, it will be minimized when the expected log probability $E_p^K = -\frac{1}{N} \sum \log p(\theta_h|\theta_d, \Theta^K)$ is minimized. We may examine the rate of change of E_p^K while decreasing the number of lobes K from a predetermined upper bound \tilde{K} to determine the optimal K —the smallest K that achieves a statistically set error tolerance. Similar to [12], we adopt the Williams’ statistical test [32] to determine this optimal K . We use critical values for 1% tolerance. Details are omitted due to limited space and readers are referred to [5, 12] for further details about the Williams’ test in general.

6. Modeling Measured BRDF Data

We evaluate the accuracy of the directional statistics BRDF model on the real-world BRDF data collected by Matusik et al.[16], which is also used in other extensive comparative studies of parametric and nonparametric BRDF models [20, 25], to provide a fair and thorough comparison. This database consists of 100 measured isotropic BRDFs sampled over $\theta_h \in [0, \frac{\pi}{2}]$, $\theta_d \in [0, \frac{\pi}{2}]$, and $\phi_d \in [0, \pi]$ ⁵, one degree apart except for θ_h where the sampling is nonlinear to achieve denser sampling around the center $\theta_h = 0$ [16]. We uniformly subsample the incident light direction represented by θ_d with 5 degrees spacing, resulting in 18 slices of each BRDF $f_r(\theta_h, \phi_d, \theta_d = \theta_d^i)$, where $\theta_d^i = \{\frac{m\pi}{36} | m = 0, \dots, 18\}$. For each slice, we fit the directional statistics BRDF model using the EM algorithm described in Section

4.

When computing the BRDF value for a given (θ_h, θ_d) pair, we linearly interpolate the nearest DSBRDF parameter values to obtain the parameter values of the DSBRDF model for that specific θ_d . This is a powerful advantage of the DSBRDF model since it automatically provides an adapted interpolation that captures the otherwise highly non-linear variation along the θ_d dimension. This is in sharp contrast to interpolation in the raw intensity space that non-parametric models would have to rely on. We empirically found that uniform sampling of θ_d with step sizes larger than 5 degrees suffices for most BRDFs. We also plan to explore parametric modeling of the DSBRDF parameter values across θ_d in the future to arrive at an even more compact representation.

We found that data points very close to $\theta_h = 0$ were unreliable for fitting the DSBRDF model, since for several extremely shiny material their intensities could be too high to reliably capture even with a high-dynamic range imaging setup. Thus, we excluded data points with θ_h less than a degree and instead added a delta term to the mixture model to compensate for the discrepancy at $\theta_h = 0$ whose height was computed after fitting the DSBRDF model.

Figure 1(a) shows some of the results of fitting 3-lobe DSBRDF models to the $\theta_d = 0$ slices of different measured BRDF data as the 1D profiles on the incident plane. The results show that the directional statistics BRDF model fits the measured data very well. Notice that the measured BRDF data exhibit various types of distributions which cannot be modeled with a simple combination of Lambertian and Torrance-Sparrow reflection models. The flexibility of individual hemi-EPDs enables us to fully capture the wide range of distribution shapes of real-world isotropic BRDFs. Although examples are omitted due to limited space, some BRDFs exhibit non-centered lobes (ones centered at the perimeter of θ_h). These are mainly caused by subsurface

⁵Again, for isotropic BRDFs, the reflected radiance values are symmetric about the incident plane.




Figure 2: Synthetic spheres rendered using the directional statistics BRDF model with parameters values estimated from measured BRDF data (nickel, specular-blue-phenolic, and orange-paint in [16]). Each row shows spheres rendered using the original measured BRDF data and those rendered using the DSBRDF model with 1 to 5 lobes from left to right, respectively. The estimated optimal number of lobes for the DSBRDF models (see Section) were 5, 4, and 2, respectively, which agree well with the visual quality of the rendered spheres.

scattering⁶, which cannot be captured with the directional statistics BRDF model. We plan to extend the model to include such non-centered lobes. In this paper, we use a simple circular mask on (θ_h, ϕ_d) with a radius proportional to θ_d to discard data points at the θ_h perimeter.

Figure 2 shows synthetic spheres rendered under an environmental illumination (St. Peter’s Basilica light probe in [7]) using the original tabulated measured BRDF data [16] as well as the fit parametric directional statistics BRDF model with 1 to 5 lobes, for different BRDFs. We implemented the rendering in `pbrt` as a new material with a new `BxDF` [22]. One can see that the rendered spheres become visually indistinguishable from the measured BRDF after the optimal number of lobes is used in the DSBRDF model, which were 5, 4, and 2, respectively. These results clearly show that the DSBRDF model captures the behavior of the BRDF extremely well with only a small number of parameter (measured BRDF requires $3 \times 90 \times 90 \times 180$ floating point values while the DSBRDF model with K lobes only needs $3 \times 2K \times 18$).

The directional statistics BRDF model automatically provides a decomposition of the BRDF into their constituent lobes. Figure 1(b) shows the 1D profiles of the three individual lobes of 3-lobe DSBRDF models fit to measured data (we are only showing the red channel). Each lobe clearly models a separate reflection component with distinct characteristics. For instance, we may interpret the first two lobes, the red and green curves, in the left most DSBRDF model as the specular spike and the specular lobe. Our DSBRDF model

reveals that even the specular spike has a certain spread in the angular domain and successfully models it, which would otherwise be extremely hard to model with other analytical specular reflection models. The third lobes (blue) in the 3rd and 4th DSBRDF models correspond to the diffuse lobes representing body reflections, which again cannot be modeled with a simple Lambertian model.

Figure 3 shows synthetic spheres each rendered with the individual constituent lobe (hemi-EPD) of a 3-lobe DSBRDF model fit to the measured data. The renderings clearly visualize the distinct reflectance characteristics of individual lobes. One can see that for a very glossy BRDF, the three lobes are dedicated to modeling specularity of different roughness (top two rows), while for materials with some matte appearance, the 3rd lobe essentially separates out the diffuse reflection that solely encodes the color of the material (bottom two rows).

Note that such decomposition of the BRDF cannot be achieved with non-parametric BRDF models. We believe this is a strong contribution that can benefit many applications in computer vision, such as reflection component separation for robust recognition, 3D reconstruction, etc.

Figure 4 shows the relative (energy-normalized) RMS errors [25] for all the 100 BRDFs in [16] using DSBRDF models with optimal numbers of lobes. The results show that the directional statistics BRDF model can model a wide-range of real-world isotropic BRDFs accurately. The overall accuracy is comparable to the non-parametric bivariate model introduced in (c.f. Fig. 3 in [25]). The significance of the DSBRDF model lies in the fact that we can achieve this accuracy with a very small number of parameters and thus with significantly smaller footprint.

Figure 5 shows how the relative RMS errors of the DS-

⁶Since the BRDF data in [16] are collected by capturing images of a sphere painted with the specific material of interest, subsurface scattering can contribute to strong intensities at the perimeter of the sphere, thus leading to sharp lobes centered around the perimeter of θ_h for large θ_d .




Figure 3: Synthetic spheres rendered with measured BRDF, 3-lobe DSBRDF model, and 1st to 3rd individual lobes of the DSBRDF model from left to right, respectively. The lobe decompositions clearly visualize the distinct reflectance characteristics of individual lobes, e.g., the color is solely encoded in the 3rd lobe for the bottom two materials indicating body reflection.




Figure 4: Relative RMS errors for all 100 BRDFs in [16] using the DSBRDF model with optimal number of lobes and synthetic spheres rendered with the DSBRDF model (first row) and measured data (second row). The large errors are mainly caused by subsurface scattering (see the left most column); otherwise the visual quality of the DSBRDF model is very high.

BRDF model and a non-parametric representation [16] with linear interpolation varies as the sampling of the measured data is reduced. The results clearly show that the DSBRDF model achieves higher accuracy than the non-parametric model that heavily relies on densely sampled measurements and interpolation among the sampled data points, even when the data is only reduced by a moderate amount. These results clearly demonstrate the importance of having an accurate low-dimensional parametric BRDF model for modeling real-world data, since sampling densities are usually sparse




Figure 5: Relative RMS errors of the DSBRDF model (solid lines) and a non-parametric representation with linear interpolation (dashed lines) for three different BRDFs (distinct colors) as the sampling of the measured data is reduced (we subsample indices of θ_h in [16] by the integer factors on the horizontal axis). The results show that the DSBRDF model achieves higher accuracy than a non-parametric representation even with moderate subsampling, demonstrating the robustness of the DSBRDF model.

in real-world applications.

7. The Space of Isotropic BRDFs

By modeling various measured BRDF data with the directional statistics BRDF model and storing all parameter values for all slices into a single vector, we obtain the coordinates of each BRDF in a low-dimensional space ($3 \times 6 \times 18$ -dimensions for 3 lobes and 18 slices) where each parameter set corresponds to each other between different BRDFs—we have a consistent decomposition of all the BRDFs. These parameters directly provide useful handles to examine and characterize the space of real-world isotropic BRDFs.

For instance, we may visualize the space by embedding the parameter values into a two-dimensional space. Figure 6 shows the results of applying ISOMAP [29] to the 3-lobe directional statistics BRDF model parameter values for $\theta_d = 0$. Although difficult to visualize in the figure, there is a dense cluster on the left which contains most purely matte BRDFs. The embedding clearly layouts a coordinated spectrum of isotropic BRDFs, especially in terms of their degrees of specularity in the horizontal direction from left to right. Such physically intuitive embedding strongly suggests that the DSBRDF model successfully characterizes the intrinsic structure of the BRDF space.

8. Conclusion

We introduced a novel parametric BRDF model based on the idea of modeling BRDFs as a set of directional statistics distributions. For this, we derived a new hemispherical distribution and defined the BRDF as a collection of 2D slices where each individual slice is modeled as a mixture of these hemispherical distributions. We showed that the novel




Figure 6: The space of isotropic BRDFs visualized in two dimensions by embedding the parameter values of 3-lobe DSBRDF models fit to measured data. Starting with the left most cluster of very matte BRDFs, the horizontal transition encodes the continuum of isotropic BRDFs with increasing specularity towards the right.

directional statistics BRDF model can accurately model a wide variety of real-world isotropic BRDFs achieving accuracy comparable to non-parametric models and also achieving higher accuracy than them with less data.

We believe that the new parametric BRDF model has strong implications for a broad range of applications. We have merely scratched the surface of its advantages by demonstrating its ability to automatically decompose measured BRDF into physically meaningful reflection components and its use for exploring the entire space of isotropic BRDFs. The model lays the foundation for deriving canonical probabilistic formulations (as we showed for estimating its parameters) for long standing physics-based vision problems, which we plan to investigate in the future.

Acknowledgement

This work was supported in part by National Science Foundation CAREER Award IIS-0746717.

References

- [1] M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Dover, 1965.
- [2] M. Ashikhmin and S. Premoze. Distribution-based BRDFs. Technical report, University of Utah, Mar. 2007.
- [3] C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- [4] J. Blinn. Models of Light Reflection for Computer Synthesized Pictures. *ACM SIGGRAPH*, 11(2):192–198, 1977.
- [5] S. Cang and D. Partridge. Determining the Number of Components in Mixture Models using Williams’ Statistical Test. In *Int’l Conf. Neural Information Processing*, pages 14–18, 2003.
- [6] R. Cook and K. Torrance. A Reflectance Model for Computer Graphics. *ACM Trans. on Graphics*, 1(1):7–24, 1982.
- [7] P. Debevec. Light Probe Gallery.
- [8] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, and W. Sarokin. Acquiring the Reflectance Field of a Human Face. In *Proc. of ACM SIGGRAPH 2000*, pages 145–156, 2000.
- [9] D. Edwards, S. Boulos, J. Johnson, P. Shirley, M. Ashikhmin, M. Stark, and C. Wyman. The Halfway Vector Disk for BRDF Modeling. *ACM Trans. on Graphics*, 25(1):1–18, 2006.
- [10] R. Fisher. Dispersion on a Sphere. *Proc. the Royal Society of London Series A* 217, pages 295–305, 1953.
- [11] A. Ghosh, S. Achutha, W. Heidrich, and M. O’Toole. BRDF Acquisition with Basis Illumination. In *Proc. of Int’l Conf. on Computer Vision*, 2007.
- [12] K. Hara, K. Nishino, and K. Ikeuchi. Mixture of Spherical Distributions for Single-View Relighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(1):25–35, Jan. 2008.
- [13] J. Koenderink and A. van Doorn. Phenomenological Description of Bidirectional Surface Reflection. *OSA Journal of Optical Society America, A*, 15(11):2903–2912, 1998.
- [14] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Computer Graphics Proceedings, ACM SIGGRAPH 97*, pages 117–126, Aug. 1997.
- [15] J. Lambert. *Photometria sive de mensura de gratibus luminis colorum et umbrae*. Augsburg, 1760.
- [16] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A Data-Driven Reflectance Model. *ACM Trans. on Graphics*, 22(3):759–769, 2003.
- [17] W. Matusik, H. Pfister, M. Brand, and L. McMillan. Efficient Isotropic BRDF Measurement. In *Proc. of Eurographics Symposium on Rendering*, 2003.
- [18] S. Nayar, K. Ikeuchi, and T. Kanade. Surface Reflection: Physical and Geometrical Perspectives. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):611–634, 1991.
- [19] S. Nayar and M. Oren. Generalization of the Lambertian Model and Implications for Machine Vision. *International Journal of Computer Vision*, 14:227–251, 1995.
- [20] A. Ngan, F. Durand, and W. Matusik. Experimental Analysis of BRDF Models. In *Eurographics Symposium on Rendering*, pages 117–226, 2005.
- [21] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. Geometric Considerations and Nomenclature for Reflectance. National Bureau of Standards (US), 1977.
- [22] M. Pharr and G. Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann, 2004.
- [23] B. Phong. Illumination for Computer Generated Pictures. *Communications of ACM*, 18(6):311–317, 1975.
- [24] R. Ramamoorthi and P. Hanrahan. A Signal-Processing Framework for Inverse Rendering. In *Computer Graphics Proceedings, ACM SIGGRAPH 01*, pages 117–128, Aug. 2001.
- [25] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive Reflectometry. In *Proc. of European Conf. on Computer Vision*, volume IV, pages 859–872, 2008.
- [26] S. Rusinkiewicz. A New Change of Variables for Efficient BRDF Representation. In *Eurographics Workshop on Rendering*, 1998.
- [27] C. Schlick. An Inexpensive BRDF Model for Physically-Based Rendering. *Computer Graphics Forum*, 13(3):233–246, 1994.
- [28] M. Stark, J. Arvo, and B. Smits. Barycentric Parameterizations for Isotropic BRDFs. *IEEE Trans. on Visualization and Computer Graphics*, 11:126–138, 2005.
- [29] J. Tenenbaum, V. Silva, and J. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290:2319–2323, 2000.
- [30] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America*, (57):1105–1114, 1967.
- [31] G. Ward. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH 92*, pages 265–272, 1992.
- [32] D. Williams. A test for differences between treatment means when several dose levels are compared with a zero dose control. *Biometrics*, 27:103–117, 1971.