

Technical Report DU-CS-10-06

Department of Computer Science, Drexel University

A Directional Statistics-based Reflectance Model for Isotropic BRDFs

Ko Nishino Stephen Lombardi

Department of Computer Science

Drexel University

3141 Chestnut Street

Philadelphia, PA 19104, USA

E-mail: 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 real-world isotropic BRDFs as mixtures of it. We derive a canonical probabilistic method for estimating the parameters, including the number of components, of this novel directional statistics BRDF (DSBRDF) model. We show that the model captures the full spectrum of real-world isotropic BRDFs with high accuracy but with a small footprint. We also demonstrate the advantages of the novel BRDF model by showing its use for reflection component separation and for exploring the space of isotropic BRDFs.

I. INTRODUCTION

The appearance of a real-world object is determined by how light rays interact with the object surface. In the most general form, this interaction can be mathematically modeled with the General Scattering Function [21] which is a function of fourteen variables in total. In reality, we may make various assumptions about the object surface that lead to significant reduction in the number of variables of this function. If we assume that the light interaction is strictly local, in other words, we are only concerned with light transport at a surface point and choose to ignore all subsurface transports, e.g., subsurface scattering in layered translucent materials, the mathematical model can be reduced to the Bidirectional Reflectance Distribution Function (BRDF) [21].

Although limited to expressing strictly local light interaction, the BRDF is arguably the most widely used description of light transport in computer vision and graphics. This is particularly true in applications where properties of the scene, such as the illumination, 3D geometric structure, and reflectance properties are inferred from images, since the pointwise description enables estimation of these quantities at each scene point. As such, devising a compact yet accurate representation of real-world BRDFs goes to the heart of a number of important applications. In particular, low-dimensional parametric BRDF models play a vital role since they enable the formulation of such representative inverse problems as model parameter estimation. Yet, deriving a low-dimensional parametric BRDF model that can accurately encode the drastic variability of real-world BRDFs remains a challenging problem.

Various parametric BRDF models, either based on purely empirical observations, such as Lambertian [15], Phong [23], [4], Schlick [28], physically-based modeling of the microscopic surface geometry, including Torrance-Sparrow [6], [30] Ward [31], Oren-Nayar [19], 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], have been introduced in the past. 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 gonireflectometers, 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], [26], [29] 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 of real material, 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 deducing a lower-dimensional non-parametric representation, they hardly result in a compact one 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 dense measurements to achieve certain accuracy. This is, in turn, naturally means that one will need more data when solving inverse problems using such models compared to parametric models.

The goal of this paper is to derive a low-dimensional parametric BRDF model that can achieve accuracy comparable to non-parametric representations. That is, we aim to accurately express the wide variety of real-world BRDFs with an analytical model consisting of a small number

of parameters. In order to make the problem tractable, we will focus on isotropic BRDFs. Our approach to this challenging, long-standing problem is based on a novel perspective of a BRDF. 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 as a mixture of it. The novel directional statistics BRDF model (DSBRDF) can encode 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 analytical expression and also derive a canonical probabilistic method for estimating its parameter values including the optimal number of mixture components. Experimental results show that the model achieves accuracy comparable to the state-of-the-art non-parametric model [26].

We believe the novel directional statistics BRDF model has direct implications in a broad range of applications. Most important, it enables the encoding of a wide variety of real-world isotropic BRDFs with a common low-dimensional analytical form. This is in sharp contrast to previous work where the appropriate model had to be chosen and combined among different BRDF models or the representations were left high-dimensional incurring burden on further analysis. We demonstrate the significance of the novel model by showing that it 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, most significant, it provides powerful means to explore and characterize the entire space of isotropic BRDFs.

II. 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 ,

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

¹Note that this is very 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 much compact yet flexible BRDF model.

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 $\int_{\Omega} f_r(\omega_i, \omega_o) d\omega_i \leq 1$.

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. This is particularly true for those materials that exhibit so-called isotropic BRDFs. As is also done to derive a bivariate non-parametric BRDF model [26], we may exploit this lower-dimensional intrinsic characteristic of real-world isotropic BRDFs by making a few realistic assumptions and coordinate change.

First, for 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 $f_r(\theta_i, \theta_o, |\phi_i - \phi_o|)$ [29]. Second, we leverage a common reparametrization known as the halfway vector representation [27]. 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 [27].

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

With the halfway representation, an isotropic distribution around the halfway vector, i.e. axially symmetric distribution about $\theta_h = 0$, becomes 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 (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 the azimuthal angles ϕ_h or ϕ_d ³. Recently, Romeiro et al. experimentally 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 their values over ϕ_d [26].

To this end, we will model real-world isotropic BRDFs with essentially two-dimensional representations as we detail next. We, however, do not simply drop the dependency on ϕ_d (or ϕ_h), but instead seek to estimate an optimal **parametric** representation for $f_r(\theta_h, \theta_d, \phi_d)$ that fits measured real-world isotropic BRDFs in the least squares sense⁴.

III. HEMISPHERICAL EXPONENTIAL POWER DISTRIBUTION

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 $\theta_d^i = \{\frac{m\pi}{2M} | m = 0, \dots, M \in \mathbb{Z}\}$. We

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

⁴Averaging is also optimal in a least-square sense but cannot be adapted to incorporate more sophisticated analysis such as robust estimator-based outlier rejection.

will later discuss how the parameters of these sampled slices can be consolidated into a smaller set of parameters through functional data analysis. Figure 2(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 encode the large variability 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 are, however, not suitable to model these BRDF slices. For instance, a widely used directional statistics distribution is the von Mises-Fisher distribution [10], a Gaussian distribution on the unit sphere

$$p(\theta_h|\theta_d, \Theta) = \frac{\kappa}{4\pi \sinh \kappa} \exp [\kappa \cos \theta_h] , \quad (3)$$

where the only parameter is $\Theta = \{\kappa\}$. κ controls the “concentration” analogous to the variance of a Gaussian in Cartesian coordinates. Although one might think that the von Mises-Fisher distribution (or other sphere-wrapped Gaussian distributions) may suffice to model isotropic BRDFs, as shown in Figure 1(a), it clearly lacks the flexibility (note again that even a Torrance-Sparrow is not a Gaussian on the sphere) to model a wide variety of real-world isotropic BRDFs. An even bigger problem is that the domain is the entire sphere, making it unsuitable for modeling a BRDF.

In order to model various BRDFs, we need a more flexible directional statistics distribution model that is naturally defined on the unit hemisphere. 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):

$$p(\theta_h|\theta_d, \Theta) = C(\Theta) (\exp [\kappa \cos^\gamma \theta_h] - 1) , \quad (4)$$

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. As illustrated in Figure 1(b), 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 radiance, and γ controls the kurtosis of the distribution, controlling 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 $p(\theta_h|\theta_d, \Theta) = C(\Theta)(e^\kappa \exp[-\frac{\kappa\gamma}{2}\theta_h^2] - 1) + O(\theta_h^n)$, the hemi-EPD subsumes a Gaussian distribution. This means that a BRDF model based on the hemi-EPD embodies the Torrance-Sparrow model which is known to achieve high-accuracy among past analytical models.

IV. 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 reflected light rays. As a result, real-world BRDFs tend to exhibit a complex distribution that cannot be modeled with a single analytical distribution 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 or 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.

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 not just for reflectance estimation but also in other inverse problems including illumination estimation and geometry reconstruction of unknown material objects. On the other hand, when the constituent BRDFs models are the same but overly simplistic, e.g., Gaussian distributions with different parameters (multiple lobes of Torrance-Sparrow specular reflection), the resulting BRDF model would lack the expressiveness to represent a wide variety of materials.

We model real-world isotropic BRDFs as 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. An isotropic BRDF can thus be modeled as a linear combination of hemi-EPDs,

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

where $\Theta = [\Theta^{(1)} \dots \Theta^{(K)}]$ and K is the number of constituent hemi-EPDs. 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

$$p(\theta_h|\theta_d, \Theta) = \sum_{k=1}^K \frac{1}{C(\Theta^{(k)})} p(\theta_h|\theta_d, \Theta^{(k)}). \quad (6)$$

to them. We refer to this novel analytical reflection model as the *isotropic directional statistics BRDF model* or in short the *isotropic DSBPDF model*. For brevity we will drop the term isotropic in the rest of the paper. 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 $\sum^K C(\Theta^{(k)})$ to normalize its total energy. This in turn means that as long as $\sum^K \frac{1}{C(\Theta^{(k)})} \leq 1$, the energy conservation law is satisfied.

We may derive a canonical algorithm for fitting this isotropic DSBPDF model to measured data by following the same principal as that of the Expectation Maximization algorithm [3].

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 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 DSBPDF model to the 2D slices of a measured BRDF data ($\{\theta_d^i | i = 1, \dots, K\}$), 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 DSBPDF model in a least square sense, i.e. estimate the parameter values that explains the asymmetry with minimum squared error.

We may further model the joint variation of the two parameters (κ, γ) across the discrete samples of the incident light direction θ_d . This results in a DSBPDF model that expresses isotropic BRDFs as mixtures of hemi-EPDs whose parameters are functions of θ_d . In particular, in Section VI-C, we detail an approach based on functional principal component analysis to extract least-square optimal, data-driven bases for expressing these functional variations of the parameter values from real-world measured data.

V. 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. For instance, a simple BRDF that only exhibits diffuse reflection may be modeled with a single lobe, while BRDFs with exotic distributions may require more than 3 lobes each representing different types of reflections ranging

from a specular spike to a glossy specular lobe and a Lambertian lobe. We can automatically determine the optimal number of lobes by running the EM algorithm 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 $E_p^K = -\frac{1}{N} \sum^N \log p(\theta_h|\theta_d, \Theta^K)$ is minimized. Here we have denoted the parameters of a K -lobe directional statistics BRDF model with Θ^K . To this end, 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.

VI. MODELING REAL-WORLD ISOTROPIC BRDFS

We evaluate the accuracy of the directional statistics BRDF model with the real-world BRDF data collected by Matusik et al.[16] that is also used in other extensive comparative studies of parametric and nonparametric BRDF models [20], [26], 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 $\theta_d^i = \{\frac{m\pi}{36} | n = 0, \dots, 18\}$. For each slice, we fit the directional statistics BRDF model using the EM algorithm described in Section IV. We found that data points very close to $\theta_h = 0$ were unreliable for fitting the DSBPDF 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 DSBPDF model.

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

A. Representing Measured Data

Figure 2(a) illustrates some of the results of fitting 3-lobe DSBRDF models to the $\theta_d = 0$ slices of different measured BRDF data, shown 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. Figure 3 shows the results of fitting the DSBRDF model to different slices of two different measured BRDF data. The results show that the novel BRDF model can accurately model the variation, most notably the increased overall intensity and the shape change of specular lobes/spikes as the incident light direction approaches grazing angles. Notice that some BRDFs (second row of Figure 3) exhibit non-centered lobes (the ones at the perimeter of θ_h). These are mainly caused by subsurface scattering⁶, which cannot be captured with the DSBRDF 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 4 shows synthetic spheres rendered under an environmental illumination (St. Peter’s Basilica light probe in [7]) using the original tabulated, measured BRDF data 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 $2 \times 3 \times K \times 18$, $(\kappa, \gamma) \times (R, G, B) \times (K - \text{lobes}) \times (\theta_d - \text{sampling})$, where K is typically 3. Again, in Section VI-C, we further reduce the size of this representation through functional principal component

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 .

analysis on the variation of the two parameters (κ, γ) across θ_d samples. Even without it, however, the DSBPDF model is much more compact representation (over a ten thousand times reduction in size) compared to a straightforward non-parametric representation.

Figure 5 shows the relative (energy-normalized) RMS errors for all the 100 BRDFs in [16] using DSBPDF 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 [26]), which achieves higher accuracy than past parametric models⁷. The significance of the DSBPDF 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 6 shows how the relative RMS errors of the DSBPDF model and a non-parametric model with linear interpolation varies as the sampling of the measured data is reduced. The results clearly show that the DSBPDF 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 in real-world applications.

B. Reflection Component Separation

The directional statistics BRDF model automatically provides a decomposition of the BRDF into their constituent lobes. Figure 2(b) shows the 1D profiles of the three individual lobes of 3-lobe DSBPDF 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 lobes, in the left most DSBPDF model as the specular spike and the specular lobe. Our DSBPDF 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 DSBPDF models correspond to the diffuse lobes representing body reflections, which again cannot be modeled with a simple Lambertian model.

⁷Again, note that we exclude data points with $\theta_h < 1^\circ$, which may explain why the errors are slightly less than that of [26].

Figure 7 shows synthetic spheres each rendered with the individual constituent lobe (hemi-EPD) of a 3-lobe DSBPDF 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 specularly 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).

Figure 8 shows the results of modeling an interesting material, polyoxymethylene plastic, with a 3-lobe DSBPDF model. The DSBPDF model accurately captures the overall characteristics of the BRDF as is evident in the precise fitting of the lobe profiles for different color channels (shown for the θ_d slice in the incident plane in the second column) and the rendered sphere under natural illumination (first column). The reflection components are clearly separated into three distinct lobes (third column) each capturing a constant component in the angular domain (green curve), a shallow specular-like lobe (blue curve), and a Gaussian-like lobe (red curve). It is, however, interesting to note that the second lobe, when rendered from the estimated DSBPDF model parameters, encodes body color and does not purely reflect the illumination color. This is contrary to the widely adopted dichromatic model in which the body color is solely encoded in the diffuse component (green curve) and the illumination color is solely encoded in the specular component (red curve). It is important to note that conventional reflection models often blindly used in computer vision, i.e., Lambertian plus Torrance-Sparrow reflection models, cannot model such phenomenon, while the DSBPDF model accurately captures the unorthodox reflection components.

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., which we would like to explore in the future.

C. The Space of Isotropic BRDFs

The directional statistics BRDF model can represent a continuum of real-world isotropic BRDFs with the same low-dimensional parametric form. This enables the analysis of the space of isotropic BRDFs, represented with measured data of 100 samples [16], with a unified expression. This ability is vital to study the space and characterize it for further use in various inverse or

forward problems – for instance to extract statistical priors to constrain ill-posed inverse problems or to realistically vary the appearance of one material to another. To validate the effectiveness of using the DSBPDF model for analyzing the space of isotropic BRDFs, we demonstrate its use for computing a low-dimensional embedding of the space.

We start the analysis by first modeling the variation of the DSBPDF parameters across different incident light directions. We view each κ and γ of a single lobe as the value of a bivariate function of θ_d sampled at 18 uniformly-spaced values. We refer to this bivariate function as the (κ, γ) -curve of a particular lobe. We choose to model these (κ, γ) -curves with quadratic B-splines with five uniform knots based on an exploratory analysis of various functional bases with leave-one-out cross validation. With this interpretation of κ and γ across θ_d , we have reduced the number of parameters of the model by one third as now the 18 samples are expressed with 6th-order B-splines.

We analyze the variability of the (κ, γ) -curves by performing bivariate functional principal component analysis (FPCA) [25] on all lobes of the 100 real-world BRDFs [16]. To properly account for the scale disparity between κ and γ , we modify the inner product between (κ, γ) -curves by weighting the κ and γ portions of the product by the reciprocal of their mean curve squared

$$\langle \xi_1, \xi_2 \rangle = \frac{1}{\|\mu^\kappa\|^2} \int \xi_1^\kappa(t) \xi_2^\kappa(t) dt + \frac{1}{\|\mu^\gamma\|^2} \int \xi_1^\gamma(t) \xi_2^\gamma(t) dt, \quad (7)$$

where $\xi = (\xi^\kappa, \xi^\gamma)^T$ is a (κ, γ) -curve. The principal functions (eigenfunctions) not only capture the major variations of κ and γ separately across θ_d , but also capture how κ and γ vary together. By using these principal functions as the bases, we arrive at a natural low-dimensional subspace representation of the space of isotropic BRDFs encoded by individual lobes (BRDFs decomposed into their reflectance components).

Figure 9 shows the projection of each lobe of all 100 BRDFs onto the first and second bivariate principal functions. Each BRDF was first represented with a DSBPDF with an optimal number of lobes and then each lobe was projected into the functional eigenspace individually. The results exemplify how the principal functions naturally characterizes the isotropic BRDF space in physically meaningful ways, with pure diffuse lobes located in the top-left of the graph while specular surfaces are scattered in the bottom-right. This, in turn, strongly indicates that the DSBPDF model successfully retains the physical properties of the BRDFs in its encoding.

We further extend this analysis to characterize the BRDF space while retaining the DSBPDF representation of each BRDF as a whole. We first establish correspondences between the lobes across different DSBPDFs that have different optimal number of lobes. We cluster each lobe into two classes based on the norm of their γ curve – diffuse and specular as γ encodes the acuteness of the lobe. The lobes in one color channel of a BRDF are ordered by their γ values, and most channels have two lobes which “straddle” this border between the two classes. We establish lobe correspondences across different BRDFs along this border. By padding any lobes which do not have correspondences with flat lobes, we arrive at a 9-lobe model for each color channel and a total of 54 curves per BRDF (3 color channels per BRDF, 9 lobes per channel, and 2 curves per lobe). We perform multivariate FPCA on these 54 curves together to obtain principal functions of the (κ, γ) -curves that describe the variability of a BRDF as a whole rather than of each constituent lobe.

Figure 10 shows the projections of each BRDF onto the first two bivariate principal functions. We note that the projections give a natural clustering of the BRDFs with pure diffuse lobes clustered very tightly in the top right of the graph, plastic-looking materials located along the right side, and metallic-looking materials spread along the top of the graph. Because the first two bivariate eigenfunctions account for about 54% of the variation of the data, we expect an even better clustering when considering more than two dimensions. Using this decomposition, nine principal components are required to represent 90% of the variability, 13 principal components to represent 95%, and 22 principal components to represent 99%.

Note that such low-dimensional embedding results in a valid characterization of the massively high-dimensional, original space of isotropic BRDFs only because the DSBPDF model successfully extracts a meaningful low-dimensional parameterization of each BRDF. Such low-dimensional characterization of the space would not be possible if we used a non-parametric BRDF model, since the implicit correspondences among different BRDFs cannot be established with just non-linear dimensionality reduction and would require additional charting.

VII. 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 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 advantage of having a common low-dimensional parametric form of real-world BRDFs is probably most significant in solving inverse problems, such as illumination estimation, material property estimation, and 3D reconstruction from images. Since the DSBPDF model is also readily a probability density function, we may reformulate many of these long-standing problems as probabilistic inference in which we can leverage strong priors on real-world BRDFs extracted by, for instance, the analysis of the space of BRDFs. We plan to report on these in the future.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] M. Abramowitz and I.A. 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.M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- [4] J.F Blinn. Models of Light Reflection for Computer Sythesized 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.L. Cook and K.E. 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 Halway Vector Disk for BRDF Modeling. *ACM Trans. on Graphics*, 25(1):1–18, 2006.
- [10] R.A. 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.J. Koenderink and A.J. van Doorn. Phenomenological Description of Bidirectional Surface Reflection. *OSA Journal of Optical Society America, A*, 15(11):2903–2912, 1998.
- [14] E.P.F. Lafortune, S-C. Foo, K.E. Torrance, and D.P. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Computer Graphics Proceedings, ACM SIGGRAPH 97*, pages 117–126, Aug. 1997.
- [15] J.H. 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.K. 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.K. 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.E. Nicodemus, J.C. Richmond, J.J. Hsia, I.W. 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.T. 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] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer, 2nd edition, June 2005.
- [26] F. Romeiro, Y. Vasilyev, and T.E. Zickler. Passive Reflectometry. In *Proc. of European Conf. on Computer Vision*, volume IV, pages 859–872, 2008.
- [27] S. Rusinkiewicz. A New Change of Variables for Efficient BRDF Representation. In *Eurographics Workshop on Rendering*, 1998.
- [28] C. Schlick. An Inexpensive BRDF Model for Physically-Based Rendering. *Computer Graphics Forum*, 13(3):233–246, 1994.
- [29] M. Stark, J. Arvo, and B. Smits. Barycentric Parameterizations for Isotropic BRDFs. *IEEE Trans. on Visualization and Computer Graphics*, 11:126–138, 2005.
- [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.J. Ward. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH 92*, pages 265–272, 1992.
- [32] D.A. Williams. A test for differences between treatment means when several dose levels are compared with a zero dose control. *Biometrics*, 27:103–117, 1971.

LIST OF FIGURES

- 1 We derive a novel directional distribution model, the Hemispherical Exponential Power Distribution (b), which has a shape parameter γ in addition to the scale parameter κ that corresponds to the concentration of a von Mises-Fisher distribution (a) (shown as 1D profiles). The hemi-EPD model can represent a wide variety of hemispherical directional distributions (corresponding colors represent different values of κ and the three distinct distributions in (b) are drawn with different values of γ for the same κ). 22
- 2 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 DSBDRF model accurately fits the measured data despite the dramatically different shapes of the measured distributions (RGB correspond to RGB color channels). Row (b): The lobes (red, green, blue lines) of a 3-lobe isotropic DSBDRF model (black solid line) fit to measured data (black dots). Each lobe clearly captures a distinct characteristic reflectance component of the BRDF (please zoom into to see the small lobes in the 3rd column BRDF). 23
- 3 The directional statistics BRDF model (solid line) fit to different slices ($\theta_d = \{10^\circ, 40^\circ, 70^\circ\}$) of two measured BRDF data (square: color-changing-paint3 and teflon in [16]). The DSBDRF model successfully captures the variation of the BRDF distributions, including the overall intensity increase and shape changes of specular lobes/spikes as θ_d is varied. 24
- 4 Synthetic spheres rendered using the directional statistics BRDF model with parameter 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 DSBDRF model with 1 to 5 lobes from left to right, respectively. The estimated optimal number of lobes for the DSBDRF models (see Section) were 5, 4, and 2, respectively, which agree well with the visual quality of the rendered spheres. 25

- 5 Relative RMS errors for all 100 BRDFs in [16] using the DSBDRF model with optimal number of lobes and synthetic spheres rendered with the DSBDRF 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 DSBDRF model is very high. 26
- 6 Relative RMS errors of the DSBDRF 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 θ_h by the integer factors on the horizontal axis). The results show that the DSBDRF model achieves higher accuracy than non-parametric models even with moderate subsampling, demonstrating fragility of non-parametric models and the robustness of the DSBDRF model. 27
- 7 Synthetic spheres rendered with measured BRDF, 3-lobe DSBDRF model, and 1st to 3rd individual lobes of the DSBDRF 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. 28
- 8 Results of modeling the BRDF of polyoxymethylene plastic with a 3-lobe DSBDRF model. First column: Synthetic spheres rendered with measured BRDF (top) and the 3-lobe DSBDRF model (bottom). Second column: The DSBDRF model (solid curve) fit to $\theta_d = 0$ slices of the measured BRDF data [16] (square) shown as the 1D profile on the incident plane for the red, green, and blue color channels. Third column: The lobes (red, green, blue curves) of a 3-lobe DSBDRF model (black solid line) fit to measured data (black dots). Fourth column: Synthetic sphere rendered with the DSBDRF model parameters for each lobe separately. Note that the second lobe (corresponding to the blue curve in the third column) encodes the body color while it exhibits specular reflection, which cannot be modeled with conventional dichromatic models. 29

- 9 By projecting each lobe onto the computed eigenfunctions, we can visualize how the eigenfunctions naturally characterize the space of isotropic BRDFs and offer physically intuitive interpretations. In this image, each sphere is a single synthetically rendered lobe of a BRDF whose x and y coordinates are the projections onto the first two eigenfunctions. The diffuse lobes are located in the top-left with specular surfaces scattered in the bottom-right indicating that the first eigenfunction roughly encodes the variability of glossiness and the second eigenfunction the variability of matteness in the BRDF space. 30
- 10 We project each BRDF (represented as 54 curves) onto the first two bivariate eigenfunctions obtained via FPCA on the (κ, γ) -curves computed by representing each BRDF, as a whole, with a DSBPDF model with optimal number of lobes. Note the natural clustering in this low-dimensional embedding of the BRDF space, with pure-diffuse surfaces located in a tight cluster in the upper right-hand corner. Metallic surfaces are scattered towards the left and plastic surfaces towards the bottom. These results suggest that the DSBPDF model provides a sound foundation for extracting physically meaningful low-dimensional bases for encoding the otherwise massively high-dimensional space of real-world BRDFs. 31




Fig. 1

WE DERIVE A NOVEL DIRECTIONAL DISTRIBUTION MODEL, THE HEMISPHERICAL EXPONENTIAL POWER DISTRIBUTION (B), WHICH HAS A SHAPE PARAMETER γ IN ADDITION TO THE SCALE PARAMETER κ THAT CORRESPONDS TO THE CONCENTRATION OF A VON MISES-FISHER DISTRIBUTION (A) (SHOWN AS 1D PROFILES). THE HEMI-EPD MODEL CAN REPRESENT A WIDE VARIETY OF HEMISPHERICAL DIRECTIONAL DISTRIBUTIONS (CORRESPONDING COLORS REPRESENT DIFFERENT VALUES OF κ AND THE THREE DISTINCT DISTRIBUTIONS IN (B) ARE DRAWN WITH DIFFERENT VALUES OF γ FOR THE SAME κ).

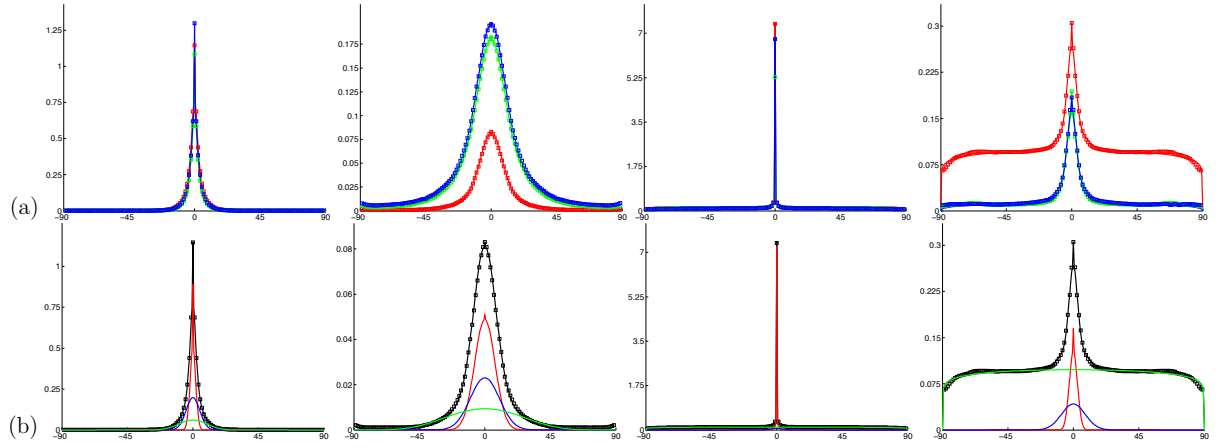


Fig. 2

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 DRAMATICALLY DIFFERENT SHAPES OF THE MEASURED DISTRIBUTIONS (RGB CORRESPOND TO RGB COLOR CHANNELS). Row (B): THE LOBES (RED, GREEN, BLUE LINES) OF A 3-LOBE ISOTROPIC DSBRDF MODEL (BLACK SOLID LINE) FIT TO MEASURED DATA (BLACK DOTS). EACH LOBE CLEARLY CAPTURES A DISTINCT CHARACTERISTIC REFLECTANCE COMPONENT OF THE BRDF (PLEASE ZOOM INTO TO SEE THE SMALL LOBES IN THE 3RD COLUMN BRDF).




Fig. 3

THE DIRECTIONAL STATISTICS BRDF MODEL (SOLID LINE) FIT TO DIFFERENT SLICES ($\theta_d = \{10^\circ, 40^\circ, 70^\circ\}$) OF TWO MEASURED BRDF DATA (SQUARE: COLOR-CHANGING-PAINT3 AND TEFLON IN [16]). THE DSBRDF MODEL SUCCESSFULLY CAPTURES THE VARIATION OF THE BRDF DISTRIBUTIONS, INCLUDING THE OVERALL INTENSITY INCREASE AND SHAPE CHANGES OF SPECULAR LOBES/SPIKES AS θ_d IS VARIED.




Fig. 4

SYNTHETIC SPHERES RENDERED USING THE DIRECTIONAL STATISTICS BRDF MODEL WITH PARAMETER 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.

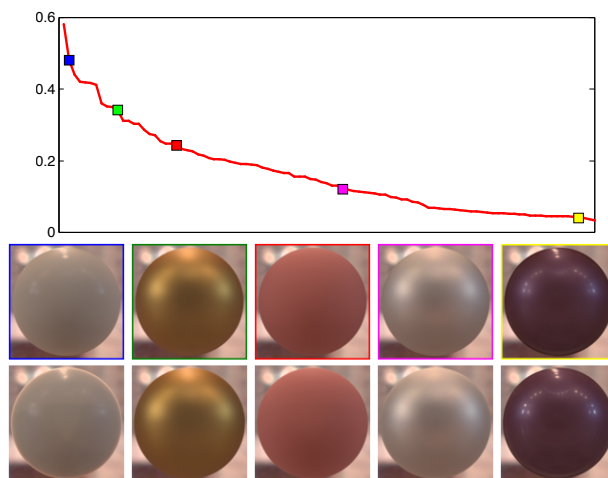


Fig. 5

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.




Fig. 6

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 θ_h BY THE INTEGER FACTORS ON THE HORIZONTAL AXIS). THE RESULTS SHOW THAT THE DSBRDF MODEL ACHIEVES HIGHER ACCURACY THAN NON-PARAMETRIC MODELS EVEN WITH MODERATE SUBSAMPLING, DEMONSTRATING FRAGILITY OF NON-PARAMETRIC MODELS AND THE ROBUSTNESS OF THE DSBRDF MODEL.




Fig. 7

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.




Fig. 8

RESULTS OF MODELING THE BRDF OF POLYOXYMETHYLENE PLASTIC WITH A 3-LOBE DSRDF MODEL. FIRST COLUMN: SYNTHETIC SPHERES RENDERED WITH MEASURED BRDF (TOP) AND THE 3-LOBE DSRDF MODEL (BOTTOM). SECOND COLUMN: THE DSRDF MODEL (SOLID CURVE) FIT TO $\theta_d = 0$ SLICES OF THE MEASURED BRDF DATA [16] (SQUARE) SHOWN AS THE 1D PROFILE ON THE INCIDENT PLANE FOR THE RED, GREEN, AND BLUE COLOR CHANNELS. THIRD COLUMN: THE LOBES (RED, GREEN, BLUE CURVES) OF A 3-LOBE DSRDF MODEL (BLACK SOLID LINE) FIT TO MEASURED DATA (BLACK DOTS). FOURTH COLUMN: SYNTHETIC SPHERE RENDERED WITH THE DSRDF MODEL PARAMETERS FOR EACH LOBE SEPARATELY. NOTE THAT THE SECOND LOBE (CORRESPONDING TO THE BLUE CURVE IN THE THIRD COLUMN) ENCODES THE BODY COLOR WHILE IT EXHIBITS SPECULAR REFLECTION, WHICH CANNOT BE MODELED WITH CONVENTIONAL DICHROMATIC MODELS.

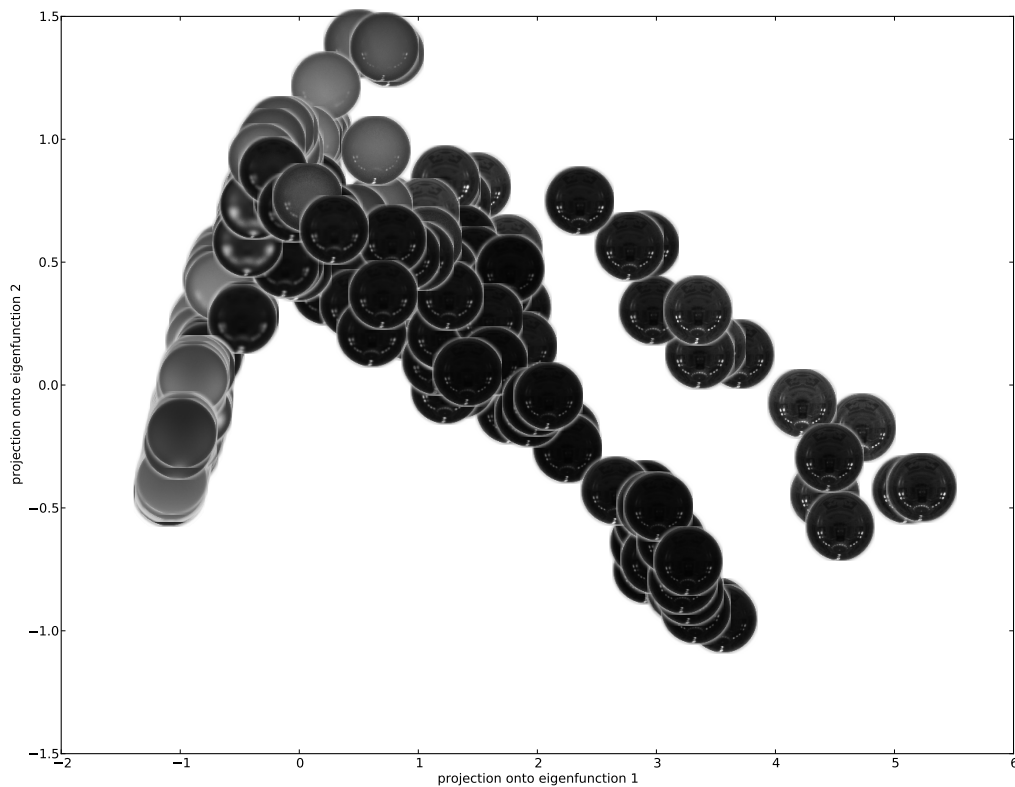


Fig. 9

BY PROJECTING EACH LOBE ONTO THE COMPUTED EIGENFUNCTIONS, WE CAN VISUALIZE HOW THE EIGENFUNCTIONS NATURALLY CHARACTERIZE THE SPACE OF ISOTROPIC BRDFs AND OFFER PHYSICALLY INTUITIVE INTERPRETATIONS. IN THIS IMAGE, EACH SPHERE IS A SINGLE SYNTHETICALLY RENDERED LOBE OF A BRDF WHOSE x AND y COORDINATES ARE THE PROJECTIONS ONTO THE FIRST TWO EIGENFUNCTIONS. THE DIFFUSE LOBES ARE LOCATED IN THE TOP-LEFT WITH SPECULAR SURFACES SCATTERED IN THE BOTTOM-RIGHT INDICATING THAT THE FIRST EIGENFUNCTION ROUGHLY ENCODES THE VARIABILITY OF GLOSSINESS AND THE SECOND EIGENFUNCTION THE VARIABILITY OF MATTENESS IN THE BRDF SPACE.

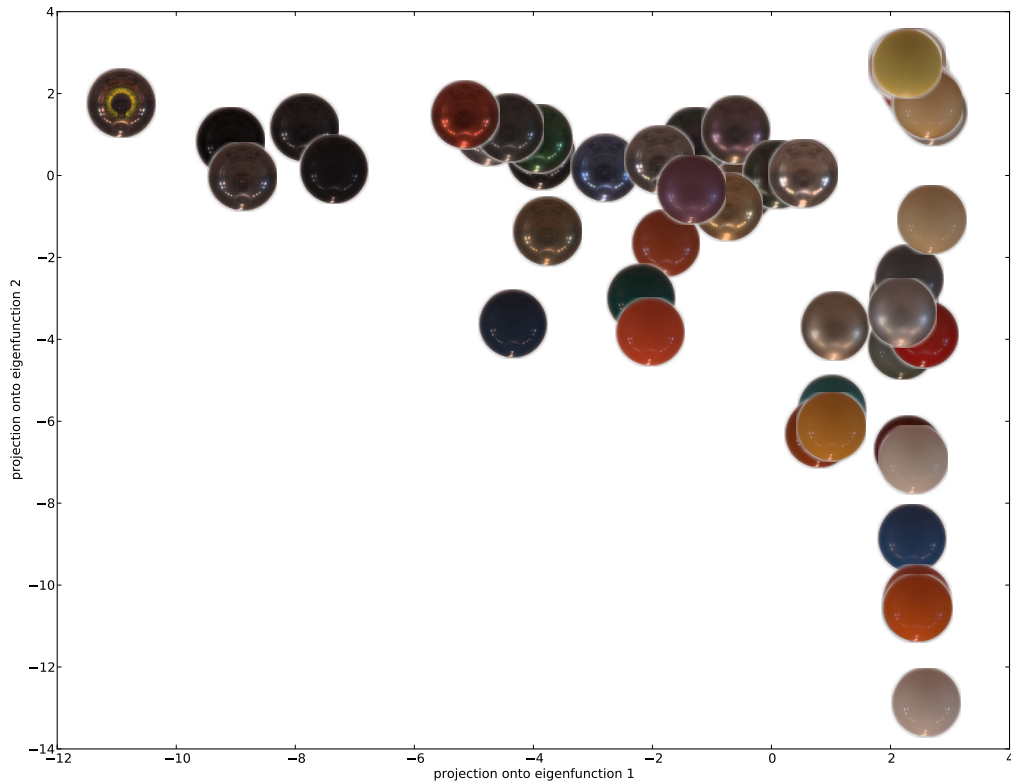


Fig. 10

WE PROJECT EACH BRDF (REPRESENTED AS 54 CURVES) ONTO THE FIRST TWO BIVARIATE EIGENFUNCTIONS OBTAINED VIA FPCA ON THE (κ, γ) -CURVES COMPUTED BY REPRESENTING EACH BRDF, AS A WHOLE, WITH A DSBRDF MODEL WITH OPTIMAL NUMBER OF LOBES. NOTE THE NATURAL CLUSTERING IN THIS LOW-DIMENSIONAL EMBEDDING OF THE BRDF SPACE, WITH PURE-DIFFUSE SURFACES LOCATED IN A TIGHT CLUSTER IN THE UPPER RIGHT-HAND CORNER. METALLIC SURFACES ARE SCATTERED TOWARDS THE LEFT AND PLASTIC SURFACES TOWARDS THE BOTTOM. THESE RESULTS SUGGEST THAT THE DSBRDF MODEL PROVIDES A SOUND FOUNDATION FOR EXTRACTING PHYSICALLY MEANINGFUL LOW-DIMENSIONAL BASES FOR ENCODING THE OTHERWISE MASSIVELY HIGH-DIMENSIONAL SPACE OF REAL-WORLD BRDFs.