

Experimental Analysis of BRDF Models

Addy Ngan, Frédo Durand,[†] and Wojciech Matusik[‡]

MIT CSAIL

MERL

Abstract

The Bidirectional Reflectance Distribution Function (BRDF) describes the appearance of a material by its interaction with light at a surface point. A variety of analytical models have been proposed to represent BRDFs. However, analysis of these models has been scarce due to the lack of high-resolution measured data. In this work we evaluate several well-known analytical models in terms of their ability to fit measured BRDFs. We use an existing high-resolution data set of a hundred isotropic materials and compute the best approximation for each analytical model. Furthermore, we have built a new setup for efficient acquisition of anisotropic BRDFs, which allows us to acquire anisotropic materials at high resolution. We have measured four samples of anisotropic materials (brushed aluminum, velvet, and two satins). Based on the numerical errors, function plots, and rendered images we provide insights into the performance of the various models. We conclude that for most isotropic materials physically-based analytic reflectance models can represent their appearance quite well. We illustrate the important difference between the two common ways of defining the specular lobe: around the mirror direction and with respect to the half-vector. Our evaluation shows that the latter gives a more accurate shape for the reflection lobe. Our analysis of anisotropic materials indicates current parametric reflectance models cannot represent their appearances faithfully in many cases. We show that using a sampled microfacet distribution computed from measurements improves the fit and qualitatively reproduces the measurements.

1. Introduction

Realistic graphics can be seen as a variant of physics that addresses light interaction. Physical sciences involve both a modeling phase where real phenomena are modeled with quantitative law, and an experimental validation phase. In realistic graphics, validation is challenging because the phenomena are complex and measuring all the scene parameters is difficult [MRC*86, RWP*95]. In this paper, we tackle an important part of the equation by focusing on material reflectance and the associated models. With the presence of a large number of BRDF models, we feel that an experimental analysis and comparison is sorely needed. In this paper, we evaluate the performance of various BRDF models by comparing their abilities to represent real measured BRDFs. We believe our analysis can serve as a guide for practitioners in their choice on BRDF models, and our extended data fitting results should also serve as starting points for using the more sophisticated models.

[†] email: addy, fredo@csail.mit.edu

[‡] email: matusik@merl.com

1.1. Related Work

BRDF research has combined skilled engineering and firm scientific principles. Empirical models respond to engineering tradeoffs between realism and efficiency, two criteria that are highlighted by the advent of programmable real-time shading. Physical models are derived from first principles and they have been compared against real data when available [CT81, HTSG91, War92].

Unfortunately, real BRDF data have long been hard to acquire and measurements were often limited in angular resolution. Recent research efforts have resulted in a wider availability of BRDF data [Cor, Nis, War92, DvGNK99, MWL*99, MPBM03, HP03, SSK03]. This provides us with a unique opportunity to experimentally analyze BRDF models and gain a practical experience on their ability to represent real BRDFs.

Understanding the tradeoffs and fitting performances of BRDFs is becoming increasingly important [SHSL97]. Inverse rendering, e.g. [FGR93, Mar98, YDMH99] and new acquisition techniques, e.g. [SWI97, LKG*01, KBMK01] typically estimate reflectance properties using a very sparse

set of observations. Even data acquired using gonireflectometers often do not contain many samples and can exhibit noise, e.g. [Bor03]. These limitations prevent the direct use of the measured data and make the fit to an analytical model a requirement. The low sampling rate and noise level make the BRDF fitting very underconstrained, and the final visual result is heavily dependent on the choice of the model. Dana et al. [DNGK97] fit a wide range of materials to BRDF models. However, their data set is limited in resolution (205 viewing/illumination combinations), and they only fit to the Oren-Nayar model [ON94] and the Koenderink representation [KvDS96] based on Zernike polynomials.

1.2. Overview

In this paper, we provide experimental analysis and insights that should guide practitioners in the field. While previous work in BRDF modeling has validated the proposed models using physical measurements, the number of materials used in the validation is typically small and often confined to the category on which the model is primarily focused. In our study we quantify and compare the performance of seven analytical BRDF models on a data set of 100 isotropic BRDFs acquired at high resolution and covering a wide range of materials. We also conduct a study on a data set of 4 common anisotropic materials measured at high resolution. The readers should refer to our [website](#)[†] for complete fitting results and instructions for raw data access.

We observe that most isotropic materials under natural illumination can be represented reasonably well using physically-based reflectance models. One of our observations is that constructions based on the mirror-angle poorly model real reflectance. In particular, we emphasize that the basis functions of the Lafortune et al. [LFTG97] model are not a good match to the shape of the primary reflection lobe of many materials (Section 4.5).

Our study indicates that some anisotropic materials can exhibit a complexity that exceeds the expressive power of current parametric BRDF models. We show that the microfacet-based BRDF generator proposed by Ashikhmin et al. [APS00] provides many more degrees of freedom and is able to reproduce these measurements qualitatively. We describe an iterative method to estimate the microfacet distribution directly from the measured BRDF (Section 5).

2. Data set and acquisition

2.1. Isotropic BRDFs

We use a data set of 100 BRDFs measured with high precision acquired by Matusik et al. [MPBM03]. The data set includes metals, plastics, painted surfaces, and fabrics. We acknowledge that our data set does not cover the full range

of materials, and effects such as retroreflection might be underrepresented. In fact, our BRDFs have no valid samples within 3 degrees from the retroreflection direction since the light source occludes the detector. The BRDFs were acquired by capturing images of a sphere sample lit by a point source from a dense set of directions, similar to Marschner et al. [MWL*99]. This image-based method allows for high angular resolution measurement since many radiance samples can be recorded in one single image. However, this measurement setup is limited to isotropic BRDFs and requires spherical samples.

2.2. Anisotropic BRDF acquisition setup

In this work, we extend the image-based BRDF acquisition setup to handle anisotropic materials (Figure 1), similar to Lu et al. [LKK00]. Spherical samples provide a two-dimensional set of normals in each image, but it is difficult to manufacture spherical samples for some materials. Marschner et al. [MWL*99] also use a cylindrical target shape for some measurements but only acquire BRDFs in the incidence plane. Using a cylindrical target any planar sample that is flexible can be wrapped on the surface without distortion. In our setup we compensate for the lost degree of freedom in the normal variations by mounting the cylinder on a precision motor and performing measurements at different tilt angles. In order to account for anisotropy the rectangular strips on the cylinder are obtained from a planar sample of the material at different orientations. Together with the degree of freedom for the light position, we are able to acquire the full 4D BRDF using a large set of two dimensional images. Figure 2 shows an example input image of an acquired velvet sample.




Figure 1: Acquisition setup for anisotropic materials. The cylinder tilt, surface normal variation, light position and strip orientation each contributes to one degree of freedom for acquiring the 4D BRDF.

The sampling density of the light position and cylinder positions can be adjusted easily as they are controlled by

[†] <http://groups.csail.mit.edu/graphics/brdf/>

precision motors. The resolution of the acquired BRDF is mainly limited by the number of material strips mounted on the cylinder. It is difficult to use strips narrower than 1cm due to the manual procedure of attaching the strips to the cylinder.




Figure 2: Target cylinder covered with velvet.

2.3. Anisotropic data

Our anisotropic data set currently includes 4 materials: velvet, two satins, and brushed aluminum. The light position is rotated at 2 degrees interval and covers most of the full circle. Positions where the light occludes the camera or is occluded by the cylinder are discarded. The motor driving the cylinder is incremented at an interval of 5 degrees, from 0 to 180. The strips on the cylinder are cut from a planar sample of the material using a laser cutter to obtain precise orientation. The strip orientation ranges from 0 to 180 degrees spaced at 9 degrees intervals. Combining this with the cylinder rotation provides a full 360 coverage. For each object/light position combination, we take 8 images with exposure time from 40 microseconds to 1 second, and combine them to form a high dynamic range image. The acquisition time is about 16 hours for each material and the raw uncompressed data size is 30GB.

2.4. Data processing

Each pixel in the input images with non-zero irradiance is considered a BRDF sample. We group all the samples into bins uniformly spaced in the 4D angular space defined by $\theta_{in}, \theta_{out}, \phi_{in}, \phi_{out}$. We set the BRDF values to the median value in each bin.

The isotropic data from Matusik et al. was originally represented using the Rusinkiewicz parameterization based on the half-vector and tabulated non-uniformly to optimize for rendering quality. To avoid bias due to the choice of parameterizations and binning intervals, we reprocess the isotropic data directly from the raw acquired images according to the same procedure described above.

We sample the isotropic and anisotropic data with intervals of 1 and 2 degrees respectively, and approximately 85% and 25% of the bins are non-empty. The processed data can be used directly for fitting by ignoring the empty bins.

3. Methodology

3.1. BRDF models

We choose seven analytical isotropic models for our analysis: Ward [War92], Blinn-Phong [Bli77], Cook-Torrance [CT81], Lafortune et al. [LFTG97], Ashikhmin-Shirley [AS00], He et al. [HTSG91], and a variant of the Ward model proposed by Duer [Due04], which fixes the normalization of the specular lobe. The formulas for the models are listed in the supplemental document on our website, and the reader is referred to the original papers for further details. These models differ in their degrees of freedom and goals, ranging from physically-based to ad-hoc models designed for computation speed. Comparing them is in a sense “unfair,” and we only focus on the numerical ability to fit measured data.

To evaluate the performance of the analytical models, we first restrict them to have only one specular lobe and the diffuse contribution to be Lambertian. The shape of the specular lobe is enforced to be identical across the color channels, and hence all the BRDF models can be expressed with six parameters for the diffuse and specular color, and a variable number of parameters for the specular lobe:

$$M = (d_r, d_g, d_b) \text{diffuse} + (s_r, s_g, s_b) \text{specular}(p_0, p_1, \dots, p_n)$$

3.2. Fitting BRDFs

A core part of our analysis is to fit measured BRDFs to analytical models. We apply constrained nonlinear optimization techniques based on Sequential Quadratic Programming (SQP) over the specular lobe parameters p_0, p_1, \dots, p_n to minimize the error metric described in the next paragraph. The diffuse and specular color (6 parameters) are computed analytically as a sub-procedure based on linear least square. Like any nonlinear optimization, the quality of the fit is dependent on a good initial guess. To make sure the optimization converges to the global minimum, we visually inspect the fitting quality of the result and if necessary restart the optimization from a different set of initial guesses.

Our fitting procedure does not take into account knowledge about the meaning of physically-based parameters for models such as Cook-Torrance or He et al. While this might result in physically-inconsistent values, it provides the best approximation of the data. Using a constrained optimization is an interesting avenue of future work.

3.3. Error Metric

The objective function of the optimization is defined as the mean squared error between the measured BRDF R , and the

target model M given parameter vector \mathbf{p} :

$$E(\mathbf{p}) = \sqrt{\frac{\sum w [R(\omega_i, \omega_o) \cos \theta_i - M(\omega_i, \omega_o; \mathbf{p}) \cos \theta_i]^2}{\sum w}} \quad (1)$$

The sum is over the non-empty bins of R , ω_i and ω_o are the incident and outgoing directions respectively, and θ_i is the elevation angle of the incident direction. The weight w is the solid angle correction term. We ignore data with incident or outgoing angle larger than 80 degrees as measurements close to extreme grazing angles are in general unreliable. The metric is essentially a L^2 metric over the hemispheres, with part of the domain removed where there is a lack of reliable measurements. The form of this metric allows us to separate the linear parameters (diffuse/specular color) from the main optimization, which greatly improves the efficiency and stability of the fitting routine.

The $\cos \theta_i$ term in the metric weighs the BRDF under the assumption of uniform incoming radiance, and in practice we observe that the optimal fit with the cosine factor produces renderings which are visually superior to the fit without it. Lafortune et al. [LFTG97] define the objective function for fitting as the BRDF difference multiplied by the cosines of both the incident and outgoing elevation angles. We find that the additional cosine factor gives results which are visually similar to ours.

We have experimented with other fitting metrics based on the logarithm or cubic root of the BRDF. As the logarithm behaves badly near zero, it imposes an arbitrary choice of scaling the BRDF, and in practice it is difficult to make a single choice that consistently performs well for all materials. For the metric based on the cubic-root, the best-fit BRDF often result in renderings with highlights that are too blurry compared to the measured data.

It is clear that choosing an appropriate metric for fitting BRDF is difficult, and our choice of the direct L2 metric is based on two reasons: i) it provides plausible fits in the rendered images for most materials in our data set, ii) the linear nature of the metric makes the minimization lower dimensional and more stable. We should emphasize that the best fit according to our metric might not always correspond to the best visual match, which is highly dependent on scene geometry/illumination. Developing a perceptually-based metric which is simple and stable for optimization would be interesting direction for future work.

4. Isotropic material analysis

4.1. Fitting quality of the models

We fit the isotropic data set of 100 materials to the seven models listed in the previous section. The fitting errors of 5 of the models are plotted in Figure 14. The errors are normalized by the maximum albedo of each BRDF and plotted in logarithmic scale. The albedo is computed by integrating

the BRDF over the outgoing hemisphere, given an incident direction. We compute the albedos of the BRDF over the range of 0 to 80 degrees incidence, and we use the maximum value to normalize the errors plotted in Figure 14. This normalization allows us to compare the residue errors for different materials while discounting the factor of relative scale. In certain cases, the absolute error can be more meaningful (e.g. darker materials are easier to approximate). The absolute errors are reported in the supplementary document available on our website.

We sort the materials according to the errors of the Lafortune fits. The errors for the Ashikhmin-Shirley and Ward-Duer models are not shown in the figure: the Ashikhmin-Shirley errors are very close to Cook-Torrance in most cases, and the Ward-Duer errors are very close to the original Ward, though consistently lower.

Figure 14 shows that He, Cook-Torrance and Ashikhmin-Shirley typically have the lowest errors. One factor of their good performance is probably their explicit modeling of the Fresnel effect. In the case of the He model, the number of degrees of freedom is also bigger. In contrast Blinn-Phong and Ward consistently yield higher errors. The errors for the Lafortune model lie in between for a majority of the materials. To the left end of the plot are mostly diffuse or mildly glossy materials like fabrics, rubber and paints, while materials close to the right end are mostly smooth metals and highly specular plastics. The normalized errors span several orders of magnitude as the shape of the BRDF is vastly different for the wide spectrum of materials.




Figure 3: Fitting the measured "green-metallic-paint" BRDF with the seven models. Clockwise from upper left: Measured, Ward, Ward-Duer, Blinn-Phong, He, Ashikhmin-Shirley, Cook-Torrance, Lafortune.

One example fit, to the measurement "green-metallic-paint" is shown in Figure 3 with spheres rendered under an environment map. The polar plots in the incidence plane for three models are also shown in Figure 4. The relative quality of the fits from the different models for this material is typical of the data set. The Ward, Ward-Duer and Blinn-Phong images are all noticeably different from the original data. The highlights near the center of the sphere are significantly




Figure 4: Polar plots of the measured "green-metallic-paint" in the incidence plane compared to the fits from the Ward, Cook-Torrance and Lafortune models. Cubic root applied.

brighter than the original, and near grazing angle the highlights are much less pronounced. Of the three models, Ward-Duer has the strongest highlight near grazing angle, but it is still noticeably dimmer than the measured data. The images from the Ashikhmin-Shirley, Cook-Torrance and He fits are very similar to each other and are mostly faithful to the direct rendering. The Lafortune fit shows an abrupt increase in intensity near grazing angle, however values closer to normal incidence are too low. In addition, the highlight is excessively blurred along the circumference of the sphere, an effect that we will look into in Section 4.5.

With a broader view over the entire data set, we find that physically-based models like Cook-Torrance and He often produce the best fits. The He model, which is known to be the most comprehensive BRDF model, does not produce noticeably superior visual results to the Cook-Torrance model. However, it should be noted that our measurement and analysis ignore polarization and spectral dependence, two important aspects modeled by He. The Lafortune model, without an explicit Fresnel factor, yields reasonable fits for materials that do not show significant intensity increase near grazing angle. In addition, glossy materials are not good candidates for the Lafortune model due to the anisotropic blurriness near grazing angle. The Blinn-Phong and Ward model both have near constant reflection power independent of incident angle, and thus they are only suitable to a subclass of the materials, including some metals and plastics. The Duer version of the Ward model with increased strength near grazing improves the fitting performance, however the improvement is limited as the grazing angle increase is not controllable.

4.2. Multiple lobes

Cook and Torrance [CT81] suggest that two or more specular lobes are necessary to model surfaces where there are multiple scales of roughness. Lafortune et al's generalized




Figure 5: The fitting errors (logarithmic scale) with one/two lobes of the Cook-Torrance and the Lafortune models.

cosine lobes [LFTG97] serve as basis functions for representing BRDFs and thus multiple lobes are expected in general. We fit the data set with the two lobes version of both the Cook-Torrance and Lafortune models. In 26 cases for Cook-Torrance and 31 cases for Lafortune, the extra lobe reduces the error by more than 25%. These typically include materials that have multiple layers of finishing (e.g. metallic paints) and also some brushed metals. The fitting errors with the Cook-Torrance and the Lafortune models with one and two lobes for the 30 most improved materials are shown in Figure 5. Figure 6 shows the comparison of the one and two lobes fit to the nickel sample.

Lafortune et al. [LFTG97] employed three lobes for the example materials. However, we found that in practice, fitting with three lobes is very unstable, and hence in the current version we omit results from the three-lobe Lafortune model. Qualitatively, we believe the addition of a third lobe would only offer marginal improvements to the shape mismatch of the main lobe.

4.3. Evaluation of different microfacet distributions

There are three common choices of the shape of the specular lobe: Gaussian, cosine power, and the Beckmann distribution function. Our study shows that the fitting quality of the Cook-Torrance model with the three different formulations are nearly identical both numerically and visually.

4.4. Fresnel effect

To study the importance of the Fresnel effect, we compare the fit quality of the Cook-Torrance model with and without the standard Fresnel term, and also with the Schlick [Sch94] approximation. While the errors without the Fresnel term are significantly higher (20% on average), the Schlick approximation does not introduce further errors compared to the original formulation.




Figure 6: Fitting the measured BRDF "nickel" with the multi-lobe Cook-Torrance and Lafortune models. Left to Right: Input, Cook-Torrance 1 lobe, Cook-Torrance 2 lobes, Lafortune 1 lobe, Lafortune 2 lobes.

4.5. Shape of the specular lobe

Our fitting results highlight the profound difference between two popular formulations of the specular lobe and their effects at near grazing angles. Most previous discussions of the behavior of reflectance when the incident angle moves toward grazing have focused on the increase in intensity of the reflection. As such, BRDF models are often only compared to measurements in the incidence plane. We find that the shape of the lobe also has important implications. The mirror angle construction employed by models like Phong and Lafortune, results in excessively blurry reflections near grazing angle. While the shape difference between these two lobe constructions is known to BRDF specialists, we believe it is important to restate it in the context of experimental evaluation, so that practitioners understand the tradeoff in models such as Lafortune that sacrifice lobe shape for computation speed.

In the original Phong model [Pho75], the specular lobe is defined around the mirror direction symmetrically ($V \cdot R$). In the Blinn variation of the Phong model and most microfacet-based models, the specular lobe is expressed with respect to the half vector H and the normal N , where H is the average vector between the incident and outgoing vector. The Lafortune model, on the other hand, employs a generalized cosine lobe[‡] similar to the original Phong lobe. While the generalized dot product gives flexibility in specifying the principal reflection direction, the lobe remains symmetric around the main reflection direction. Hence, the Lafortune lobe is similar to the $V \cdot R$ lobe where R can be seen as a generalized *mirror* direction.

Since all analytical models we have considered are functions of either $H \cdot N$ or $V \cdot R$, we focus the rest of our discussion on isocontours of these dot products in the space of directions. We illustrate the difference between the two representations in Figure 7. Consider a given incident direction

[‡] The generalized cosine lobe:

$$f_r(\mathbf{u}, \mathbf{v}) = \rho_s C_s(\mathbf{u}) [\mathbf{u}' \cdot \mathbf{v}]^n, \mathbf{u}' = \text{normalize}(C_x u_x, C_y u_y, C_z u_z) \quad (2)$$

where \mathbf{u} and \mathbf{v} are the incident and outgoing angle respectively, and $C_s(\mathbf{u})$ is a scalar factor independent of \mathbf{v} .




Figure 7: The $V \cdot R$ lobe compared with the $H \cdot N$ lobe remapped to the outgoing directions.

L . The blue circle around the normal N denotes the cone of directions where the dot product with N is constant. One of the half-vectors H in this set is shown. We can map H to its corresponding outgoing direction V_H , following a great circle on the sphere of directions. This can be done for the entire set of half-vectors in the blue circle, and we show the corresponding set of outgoing directions in green around the mirror vector R . Comparing it to the red circle, representing the cone of directions V with constant $V \cdot R$, the remapped contour is asymmetric and is narrower in the direction perpendicular to the incident plane. By similar construction, we can visualize that this contour is more circular when L is near N , and becomes narrower when L moves towards grazing angle. The $V \cdot R$ lobe is, however, always circular around the mirror direction R .

A plot of the measured BRDF of PVC plastic at 55° incident angle is shown in Figure 8. It displays a similar asymmetry to the remapped contour associated with the $H \cdot N$ lobe. This asymmetry is observed in the data set for most materials with an apparent specular lobe. Such behavior can be explained by the microfacet theory, e.g. [CT81]. The theory models the surface as a collection of microfacets, with normals oriented according to some distribution function. Assuming that the microfacets are perfect mirrors, the reflectance from L to V is then proportional to the fraction




Figure 8: Side and top-down view of the measured "PVC" BRDF, at 55° incidence. (Cubic root is applied to the BRDF for visualization.) Its specular lobe exhibits a similar asymmetry to the $H \cdot N$ lobe.

of microfacets with normals aligned with the half-vector H , with some adjustments for shadowing and masking. Thus, the reflectance function for a microfacet model is naturally defined with respect to H . For isotropic models, the distribution of normals is radially symmetric; therefore, a cone of half-vector H around N should have constant reflectance value if the shadowing and masking terms are not considered. Hence, it is reasonable to define reflectance as a function of $H \cdot N$ under this theory. On the contrary, the $V \cdot R$ contour is not consistent with any distribution of normals, as a circular cone of vectors around different R s with the same spread would map to a different set of half-vectors around N .

The difference between the two formulations explains the anisotropic blurring for the Lafortune fits in Figure 6 and Figure 3. Compared to a half-vector based lobe, the symmetric mirror lobe does not get narrow at grazing angles, resulting in blurring in the direction orthogonal to the incident plane, which is along the circumference for the case of a sphere.

We have shown that the $H \cdot N$ representation is a more accurate way to model specular reflections. This is in strong opposition to some beliefs that the half-vector representation is just a more convenient way of interpreting the specular lobe. We believe this has been overlooked in the past mainly due to the fact that the two representations are trivially related in the incidence plane.

Note that this observation is different from the popular re-parameterization proposed by Rusinkiewicz [Rus98] that also emphasizes the importance of the half vector. In his work, he shows that the parameterization based on half-angle is a more compact representation in the context of compression, compared to the standard BRDF parameterization based on the incident and outgoing directions. In our work, we compare the *shape* of the reflection lobe defined by the half vector to the one defined by the *mirror direction*.

Stark et al. [SAS05] show that Ward and Lafortune can both be mathematically reduced to bivariate functions, and thus they are inherently unable to represent general isotropic

BRDF of three variables. They went on to introduce a new set of 2D coordinates which can be considered as a hybrid of the two families of parameterization. They showed that their parameterization produces superior approximation to two measured materials when the BRDFs are projected onto the 2D domain. Further investigating the performance of Stark et al.'s proposed model and variants based on the hybrid parameterization would be an interesting avenue for future work.

5. Anisotropic material analysis

In this section we describe our first step in analyzing anisotropic reflectance models and how they are able to represent measured anisotropic materials. As described in Section 2.2, we have acquired high resolution samples for four different highly anisotropic materials.




Figure 9: Brushed Aluminum: One measurement photograph (upper left), reconstruction using Ward model ($\alpha_x = 0.04$, $\alpha_y = 0.135$) (upper right), reconstruction using Poulin-Fournier model ($n = 3000$, $d = 0.18$, $h = 0.5$) (lower left), reconstruction using Ashikhmin model with sampled microfacet distribution (lower right).

In our data set two materials (brushed aluminum and yellow satin) are suitable for analytic modeling. We fit two different anisotropic analytic models to these measurements – anisotropic Ward and Poulin-Fournier models, using the same fitting procedure as for the isotropic materials. We present the visual comparison of the fit for the brushed aluminum to the original measurements in Figure 9. As can be seen, these models can qualitatively represent brushed aluminum relatively faithfully. The yellow satin (not shown) has blurrier highlights, but the anisotropy is qualitatively similar to the brushed aluminum. As a result it can also be modeled reasonably well with the analytical models.

5.1. More Complex Anisotropy

Some fabrics that exhibit anisotropic reflectance properties can have much more complicated micro-geometry since the way threads are woven together can be potentially much more complex than the simple structure of a brushed metal. We observe that the BRDFs of the purple satin and red velvet far exceed the expressive power of simple analytical models (e.g., the purple satin sample in Figure 10 shows two off-specular peaks). This is why we turn to the microfacet-based BRDF generator presented by Ashikhmin et al. [APS00] since it offers many more degrees of freedom by replacing the analytical expression for the microfacet distribution by an arbitrary tabulated representation. In their paper, it is shown how to compute a BRDF given a microfacet distribution. In order to use their model to approximate our measured data, we need to solve the inverse problem and estimate the microfacet distribution from the BRDF.

Computing a microfacet distribution from a BRDF We first recall the expression for the Ashikhmin et al. model:

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{p(\mathbf{h}) \langle \mathbf{h} \cdot \mathbf{n} \rangle F(\mathbf{k} \cdot \mathbf{h})}{4g(\mathbf{k}_1)g(\mathbf{k}_2)} \quad (3)$$

where $p(\mathbf{h})$ is the microfacet distribution; $g(\mathbf{k}_1)$ and $g(\mathbf{k}_2)$ are shadowing and masking terms; $F(\mathbf{k}\mathbf{h})$ is the Fresnel term; and $\langle \rangle$ denotes the average over the distribution p .

We can rewrite Equation 3 in the following way:

$$p(\mathbf{h}) = \rho(\mathbf{k}_1, \mathbf{k}_2) \frac{4g(\mathbf{k}_1)g(\mathbf{k}_2)}{\langle \mathbf{h} \cdot \mathbf{n} \rangle F(\mathbf{k} \cdot \mathbf{h})} \quad (4)$$

If we ignore shadowing and masking, the BRDF is proportional to the microfacet distribution. We can then obtain a least-squares fit by averaging the estimate of $p(\mathbf{h})$ for all the pairs $\mathbf{k}_1, \mathbf{k}_2$ corresponding to the same \mathbf{h} . In practice, we guess the Fresnel term through trial and error; however, the Fresnel term can be also computed through optimization in the outermost loop.

The main difficulty, however, is the shadowing/masking function g which depends on the distribution $p(\mathbf{h})$ itself. Fortunately, we observe that g is extremely smooth [APS00] and we have found that an iterative approach works very well. We start with a constant g , and we alternate the estimation of $p(\mathbf{h})$ using the current g and the computation of g from $p(\mathbf{h})$. In practice, about five or less iterations are sufficient.

Observation on measured data We have used the above method to compute a microfacet distribution for our four measurements of anisotropic materials. Figure 9 and Figure 10 show comparisons between photographs of the measurement cylinder and renderings using the estimated microfacet distribution and the Ashikhmin et al. microfacet-based BRDF model. The renderings show good qualitative agreement with the photographs and the shape of the highlights is

well-preserved, although their sharpness is not reproduced exactly.




Figure 10: Left: Purple satin - one input photograph. Right: - reconstruction using sampled microfacet distribution.

The computed microfacet distributions also match the microstructure of the materials. Compare the spherical plot of the microfacet density in Figure 11 with a macro photograph of the corresponding sample of brushed aluminium. The density is high around a vertical great circle, which is the type of normal distribution expected from cylindrical micro-geometry.




Figure 11: Left: Brushed aluminium macro photograph. Right: deduced microfacet distribution (log plot).



Figure 12: Left: Purple satin macro photograph with sketched double cones. Right: deduced microfacet distribution (log plot).

The purple satin sample has an intriguing distribution: a large density around two great circles that are symmetrically slightly tilted from the vertical direction (Figure 12). This type of distribution can be caused by symmetric cone pairs. When we examine the macro photograph (Figure 12) we note that the longest visible segments of thread are high

in their middle and have a symmetrical slope that might be similar to a cone, as sketched on the right part of the image.




Figure 13: Left: Red velvet macro photograph. Right: deduced microfacet distribution (log plot).

The red velvet sample exhibits a secondary lobe in addition to the main lobe around the normal. This is mainly due to the dominant combing direction of the velvet sample. This shows an important type of anisotropy that cannot be modeled by the standard analytical models. We should also note that we are lacking measurements near extreme grazing angles, and therefore we may have missed important reflectance properties of the velvet in those directions. In addition, it is difficult to enforce the exact combing direction uniformly on the entire piece of velvet. Small deviations can be tolerated, however, as the average BRDF of velvet is relatively low-frequency.

6. Discussion

In this work we have analyzed the performance of several analytical BRDF models by their ability to fit real BRDFs. Our experimental results suggest that using a single specular lobe, the Cook-Torrance, Ashikhmin-Shirley and the He models perform well for most of the 100 isotropic BRDFs in the data set. Moreover, our results show that for a number of materials, the fit quality is much improved when an additional lobe is employed. This is mostly attributed to the multiple layer finish on these materials.

Our study also illustrates that the difference between the two ways of modeling the specular lobes is profound. The $H \cdot N$ formulation naturally corresponds to the microfacet-based theory, and its viability is confirmed by the data set. Meanwhile, models based on the $V \cdot R$ lobes are visually inaccurate in a rendered scene. The fitting results of these models are systematically inferior to their counterparts. The Lafortune model, though highly expressive with its basis of generalized cosine lobes, cannot accurately model real BRDFs with a single lobe. While the computation for each generalized cosine lobe is cheaper relative to a Cook-Torrance lobe, fitting more than a few of these lobes to measured data is unstable.

We have also built a new setup for high resolution measurement of anisotropic BRDFs using material strips on a rotating cylinder. We have used it to measure four typical anisotropic materials and have shown that only two of them can be modeled by parametric models. Our fabric measurements exhibit a complexity that cannot be accounted for by the elliptical functions used to introduce anisotropy in these models. Furthermore, we have shown how to estimate a tabulated microfacet distribution from these measurements, which yields good results with Ashikhmin et al.'s microfacet-based BRDF generator. This model draws its expressive power from a general representation of microfacet distributions.

There are several directions for future work. First, we expect that a new set of basis functions, similar to the generalized cosine lobes, but expressed with the half-vector, would yield improved fitting results with fewer number of lobes. Stark et al. [SAS05]'s barycentric coordinates are also possible candidates. These basis functions should be reasonably inexpensive to compute and be flexible enough to represent complex reflectance phenomena, e.g., retro-reflection. Second, while our metric for BRDF fitting gives reasonable results in most cases, it does not directly correspond to perceptually important differences. A perceptually-based metric would improve the visual quality of the fits [PFG00, RWP*95], especially when the fit is relatively poor and different choices of approximation tradeoffs are present.

This paper has also highlighted that anisotropic materials are difficult to model and acquire, and that much research effort is needed to improve the state of the art. The tabulated representation of microfacets is powerful but it requires high resolution measurements. In addition, it is not easy for an artist to edit these distributions. While the recovered distributions are more complex than elliptical shapes, they exhibit enough smoothness to leave hope for expressive and easy-to-manipulate parametric models. Finally, it is our belief that the most critical issue in reproducing real material is the acquisition itself. Our setup has allowed us to measure high quality BRDFs, but it occupies an entire room and it takes several hours to acquire one BRDF. We believe that the development of a portable and easy-to-use reflectance scanner is a major challenge.

Acknowledgement We thank Eric Chan, Jan Kautz, Jaakko Lehtinen, Daniel Vlasic and the anonymous reviewers for valuable feedback on this paper. This work was supported by an NSF CAREER award 0447561 Transient Signal Processing for Realistic Imagery, an NSF CISE Research Infrastructure Award (EIA9802220) and the Singapore-MIT Alliance.

References

- [APS00] ASHIKMIN M., PREMOŽE S., SHIRLEY P.: A microfacet-based BRDF generator. In *Proc. of SIGGRAPH 2000* (2000), ACM Press/Addison-Wesley Publishing Co. 2, 8

- [AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic phong brdf model. *J. Graph. Tools* 5, 2 (2000), 25–32. 3
- [Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. In *Proceedings of SIGGRAPH 1977* (1977), ACM Press, pp. 192–198. 3
- [Bor03] BORSHUKOV G.: Measured BRDF in film production-Realistic cloth appearance for "The Matrix Reloaded". *SIGGRAPH 2003 Sketches and Applications Program* (2003). 2
- [Cor] Cornell Light Measurement Laboratory: <http://www.graphics.cornell.edu/research/measure/>. 1
- [CT81] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. In *Proceedings of SIGGRAPH 1981* (1981), ACM Press, pp. 307–316. 1, 3, 5, 6
- [DNGK97] DANA K. J., NAYAR S. K., GINNEKEN B. V., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)* (1997), IEEE Computer Society, p. 151. 2
- [Due04] DUER A.: On the Ward model for global illumination. Submitted for review, Sept. 2004. 3
- [DvGNK99] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1 (1999). 1
- [FGR93] FOURNIER A., GUNAWAN A. S., ROMANZIN C.: Common illumination between real and computer generated scenes. In *Graphics Interface '93* (May 1993). 1
- [HP03] HAN J. Y., PERLIN K.: Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.* 22, 3 (2003), 741–748. 1
- [HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Proceedings of SIGGRAPH 1991* (1991), ACM Press, pp. 175–186. 1, 3
- [KBMK01] KOUDELKA M., BELHUMEUR P., MAGDA S., KRIEGMAN D.: Image-based modeling and rendering of surfaces with arbitrary BRDFs. In *Proc. of Computer Vision and Pattern Recognition* (2001). 1
- [KvDS96] KOENDERINK J. J., VAN DOORN A. J., STAVRIDIS M.: Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Proc. of the 4th European Conference on Computer Vision-Volume II* (1996), Springer-Verlag, pp. 28–39. 2
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proc. of SIGGRAPH 1997* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 117–126. 2, 3, 4, 5
- [LKG*01] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering* (June 2001), pp. 103–114. 1
- [LKK00] LU R., KAPPERS A., KOENDERINK J.: Optical properties (Bidirectional Reflectance Distribution Functions) of shot fabric. *Applied Optics* 39, 31 (Nov. 2000), 5785–5795. 2
- [Mar98] MARSCHNER S.: *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, 1998. 1
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (2003), 759–769. 1, 2
- [MRC*86] MEYER G. W., RUSHMEIER H. E., COHEN M. F., GREENBERG D. P., TORRANCE K. E.: An experimental evaluation of computer graphics imagery. *ACM Trans. Graph.* 5, 1 (1986), 30–50. 1
- [MWL*99] MARSCHNER S., WESTIN S., LAFORTUNE E., TORRANCE K., GREENBERG D.: Image-based BRDF measurement including human skin. In *Rendering Techniques '99* (1999), pp. 139–152. 1, 2
- [Nis] NIST reference reflectometer: STARR facility. <http://physics.nist.gov/>. 1
- [ON94] OREN M., NAYAR S. K.: Generalization of lambert's reflectance model. In *Proc. of SIGGRAPH 1994* (1994), ACM Press, pp. 239–246. 2
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proc. of SIGGRAPH 2000* (2000), ACM Press/Addison-Wesley Publishing Co., pp. 55–64. 9
- [Pho75] PHONG B. T.: Illumination for computer generated pictures. *Communications of ACM* 18, 6 (1975), 311–317. 6
- [Rus98] RUSINKIEWICZ S. M.: New change of variables for efficient BRDF representation. In *Rendering Techniques '98* (1998), Drettakis G., Max N., (Eds.), Eurographics, Springer-Verlag Wien New York, pp. 11–22. 7
- [RWP*95] RUSHMEIER H., WARD G., PIATKO C., SANDERS P., RUST B.: Comparing real and synthetic images: Some ideas about metrics. In *Eurographics Rendering Workshop 1995* (1995). 1, 9
- [SAS05] STARK M. M., ARVO J., SMITS B.: Barycentric parameterizations for isotropic brdfs. *IEEE Transactions on Visualization and Computer Graphics* 11, 2 (2005), 126–138. 7, 9
- [Sch94] SCHLICK C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. 5
- [SHSL97] SHIRLEY P., HU H., SMITS B., LAFORTUNE E.: A practitioners' assessment of light reflection models. In *Proceedings of Pacific Graphics 97* (1997). 1
- [SSK03] SATTTLER M., SARLETTE R., KLEIN R.: Efficient and realistic visualization of cloth. In *Proceedings of the 14th Eurographics workshop on Rendering* (2003), Eurographics Association, pp. 167–177. 1
- [SWI97] SATO Y., WHEELER M. D., IKEUCHI K.: Object shape and reflectance modeling from observation. In *Proc. of SIGGRAPH 1997* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 379–387. 1
- [War92] WARD G.: Measuring and modeling anisotropic reflection. In *Proc. of SIGGRAPH 1992* (1992), ACM Press. 1, 3
- [YDMH99] YU Y., DEBEVEC P., MALIK J., HAWKINS T.: Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Proc. of SIGGRAPH 1999* (1999), ACM Press/Addison-Wesley Publishing Co. 1




Figure 14: The normalized fitting errors (logarithmic scale) of five analytic models to our isotropic data set of 100 BRDFs. The BRDFs are sorted in the errors of the Lafortune model (Red) for visualization purpose.