

Distribution-based BRDFs

Michael Ashikhmin and Simon Premože




Figure 1: A comparison between various measured BRDFs and their approximations using distribution-based BRDFs. Left: measured BRDFs. Right: distribution-based BRDF approximation. The materials used are: BMW 339 paint, Ford F8 paint, silver paint, blue paint, pearl paint, fabric, leather, nickel, cacao.

Abstract

The bidirectional reflectance distribution function (BRDF) formalism is commonly used in computer graphics to represent surface reflection properties. Although many BRDF models have been proposed, most do not possess at least some of the desirable practical properties. In this paper we present a simple and flexible model which satisfies many of these requirements. We show that the proposed model provides a good approximation for many real world materials, obeys basic physical restrictions, allows straightforward hardware implementation and provides for efficient sampling in a Monte-Carlo rendering system. A procedure to fit the model to BRDF measurement data is presented which suggests a simplified way of measuring surface reflection.

Keywords: BRDF, reflection measurement, rendering, importance sampling, microfacet models

1 Introduction

For a material to be perceived correctly in a computer graphics image, the main characteristics of its interaction with incoming light should be faithfully reproduced. Although many physical processes contribute to this interaction, surface reflection has rightfully received the most attention from computer graphics community since it is typically visually the most prominent effect. Physically-based rendering systems describe reflection behavior using the bidirectional reflectance distribution function (BRDF) $\rho(\mathbf{k}_1, \mathbf{k}_2)$ [NICODEMUS et al. 1977] which describes what fraction of light coming from a given direction \mathbf{k}_1 is reflected into some other direction

\mathbf{k}_2 . (All vector quantities appear in **bold** throughout the paper.) Much research has concentrated on direct measurement of this four-dimensional function or developing suitable mathematical models.

To reproduce reflectance properties of real-world materials most faithfully, several BRDF measurement systems for computer graphics have been designed. Some are based on the classical gonio-reflectometer approach [Greenberg et al. 1997; Integra] while others use recently more popular image-based techniques [Ward 1992; Dana et al. 1999; Marschner et al. 1999]. The quality and the amount of reflection data obtained with such systems has improved significantly in recent years [Matusik et al. 2003a; Matusik et al. 2003b; Ngan et al. 2005; Günther et al. 2005; Marschner et al. 2005]. However, it remains very difficult to obtain accurate and reliable data, especially in areas near grazing angles, and near the directions of retroreflection and strong specular highlights. The measurement process itself remains very time-consuming and laborious. Significant investment is typically necessary to build a reliable measurement system.

A variety of parametric BRDF models has been proposed ranging from very simple [Phong 1975; Blinn 1977] to quite sophisticated [He et al. 1991; Stam 1999]. As an alternative to measured BRDFs, parameters of such models can be adjusted for the best visual match with real materials in what is often a rather tedious process. Somewhat surprisingly, results obtained in this manner with even some of the simplest models are often sufficient to withstand the scrutiny of the most demanding visual effects applications. We believe that this is due to two main reasons. First, the human visual system generally seems to be not very sensitive to the details of material appearance as long as its main features (such as overall color and the width of specular highlight) are reasonably accurately reproduced. Supporting this proposition, it was recently demonstrated [Matusik et al. 2003a] that many common BRDFs can be projected into a relatively low dimensional space as well as represented by a linear combination of as few as 100 BRDFs [Matusik et al. 2003b] without noticeable degradation in perceived material properties. Second, realistic illumination conditions, complex scene geometry and textures are all abundant in modern computer graphics and can be used to compensate for reflection model deficiencies. It is known that natural lighting makes it easier for humans to recognize materials [Fleming et al. 2003], but the effect of sig-

nificantly more subtle differences due to imperfect BRDFs has not been formally studied. It is, however, natural to assume that most complex environments further mask any inaccuracies in the surface reflection representation (see Section 6). In some situations, e.g. special effects for film industry, it is even possible to control view-point and to customize lighting to create more convincing imagery.

Although generally accepted in the BRDF research community, these considerations have not yet led to a sufficiently simple and flexible process of handling surface reflection in computer graphics. This is due to the multiple characteristics desirable in a useful BRDF representation, some of which are difficult to enforce and even contradictory. Specifically, if possible, one would like such a representation to:

- represent a significant number of real-world materials with sufficient accuracy for visual applications;
- use measured data and allow acquisition of the necessary information for existing materials quickly and easily;
- be able to model new materials from scratch, i.e., not rely exclusively on measured data;
- respect basic physical properties of non-negativity, reciprocity and energy conservation;
- allow efficient sampling in a Monte-Carlo rendering system;
- allow straightforward hardware implementation;
- be compact;
- have at least a semi-intuitive interpretation and be simple to use for non-BRDF experts.

Significant research on surface reflection modeling and measurement exist and excellent practical models are currently available. For a detailed review of the state of the art in BRDF research prior to 2001, the reader is referred to Siggraph course notes [Ashikhmin et al. 2001]. We believe, however, that no existing representation offers all the properties outlined above simultaneously. Using measured [Marschner et al. 1999] or simulated [Westin et al. 1992] data directly does not offer a compact solution. Recent systems based on decomposition of acquired data are quite complex. They typically target specific properties such as hardware implementation [Kautz and McCool 1999; McCool et al. 2001] or sampling efficiency [Lawrence et al. 2004] making sacrifices in other areas. With the exception of the Matusik et al. [2003a] system, modeling of novel materials is also difficult. Edwards et al. [2006] recently proposed a different domain and procedures for BRDF data fitting, also compromising on some desirable properties, in this case reciprocity. Another interesting parameterization of BRDF domain [Stark et al. 2005] was also shown to improve fitting results for several materials. In both cases, fitting procedures remain rather elaborate.

While simple Phong-like models [Ward 1992; Ashikhmin and Shirley 2000] and early microfacet-based models [Torrance and Sparrow 1967; Cook and Torrance 1981] can be successful for many materials, they do not offer sufficient flexibility. This drawback is shared by most other models which, in many cases, are difficult to fit to experimental data, are too complex to be implemented in hardware, or lack efficient sampling schemes. This includes wave physics-based models [He et al. 1991; Stam 1999] as well as those specialized for relatively narrow, although important, classes of materials [Oren and Nayar 1994; Neumann et al. 1999]. Lafortune et al. [1997] presented a technique which fits a sum of generalized Phong lobes to measured data. While successful for




Figure 2: *Geometry of reflection. Halfvector \mathbf{h} has spherical coordinates (θ, ϕ) and lies in the plane defined by incident vector \mathbf{k}_1 and outgoing vector \mathbf{k}_2 .*

some materials, the method often fails to match more complex appearance and the non-linear fitting procedure can become unstable if more than two or three lobes are involved.

Ashikhmin et al. [2000] proposed a more flexible analytic model which can successfully approximate some materials others can not [Ngan et al. 2005]. However, its rather complex mathematical form often makes it practically inconvenient. For example, a complex iterative fitting procedure is required to handle measured data. We believe that a promising new representation can be obtained by combining the flexibility of this model with the simplicity of the A&S model which, in addition to some desirable practical properties, was also demonstrated to perform well in fitting experimental data [Ngan et al. 2005] even though it was not originally intended for this purpose. We also believe that procedures currently used for fitting BRDF models to measured data are too complex, especially given limited reliability of underlying data and the absence of robust model difference metrics. We therefore propose a much simpler procedure based on our new model. Moreover, we describe how all necessary data for our approach can be obtained without resorting to inherently complex measurement of full BRDF. We call the proposed representation distribution-based BRDFs (d-BRDFs).

Throughout the paper, the emphasis is put on describing very simple techniques with results at least comparable to current state of the art which often requires much more complex procedures. Many obvious extensions of our methods, such as using additional terms in the model or resorting to strict numerical fitting of some parameters instead of adjustments based on visual feedback, exist. In most our examples, however, we found that only relatively small improvement over demonstrated results can be obtained without resorting to significantly heavier mathematical machinery and whether the extra complexity is justified in practice depends on the needs of specific application.

We present model details in Section 2. Section 3 describes a procedure for handling measured data. In Section 4 we describe a hardware implementation of the proposed model. We then demonstrate the efficiency of a Monte-Carlo sampling scheme based on the new model in Section 5. We conclude with a discussion of our approach including some of its limitations.

2 Distribution-based Reflection Model

We choose microfacet theory [Torrance and Sparrow 1967; Cook and Torrance 1981] as the general working framework. This the-




Figure 3: C_0 function $\max((\mathbf{k}_1\mathbf{n}), (\mathbf{k}_2\mathbf{n}))$ used in A&S model (left) and close in shape C_∞ bilinear patch $(\mathbf{k}_1\mathbf{n}) + (\mathbf{k}_2\mathbf{n}) - (\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})$, both having the same values on the boundary of the domain.

ory models the surface as a large number of small facets, each acting as a Fresnel mirror for incoming light according to the laws of geometric optics. Figure 2 shows the relevant geometry. We use notation (\mathbf{ab}) for the scalar (dot) product of vectors \mathbf{a} and \mathbf{b} . All vectors used are normalized so the scalar product is the cosine of the angle between the two vectors. For a facet to transfer light from the incoming direction \mathbf{k}_1 to the outgoing direction \mathbf{k}_2 , its normal direction must coincide with the half-vector $\mathbf{h} = \text{normalize}(\mathbf{k}_1 + \mathbf{k}_2)$. Correspondingly, the distribution of normal directions of these facets $p(\mathbf{h})$ combined with the Fresnel term $F((\mathbf{kh}))$ and shadowing/masking effects on the surface are responsible for surface reflection. Note that in cases where either \mathbf{k}_1 or \mathbf{k}_2 can be used in our expressions, we will drop the subscript. The theory was recently formulated for general distributions of microfacet normals by Ashikhmin et al. [2000] who noted that the shape of the reflection lobe depends mainly on this distribution function. We take this observation to the limit and provide a simple expression for the BRDF where we directly use only $p(\mathbf{h})$, standard Fresnel fraction, and an overall scaling constant as characteristics of the surface.

We start from the anisotropic Phong model (also known as A&S model) [Ashikhmin and Shirley 2000]. We first note that by rearranging the terms, we can write its specular part as:

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{c_1 p_{\text{Phong}}(\mathbf{h}) F((\mathbf{kh}))}{(\mathbf{kh}) \max((\mathbf{k}_1\mathbf{n}), (\mathbf{k}_2\mathbf{n}))} \quad (1)$$

where $p_{\text{Phong}}(\mathbf{h})$ is anisotropic Phong function normalized so that it integrates to one over hemisphere of \mathbf{h} directions and c_1 is an RGB scaling constant. One can see that in this form, the empirical A&S model can be interpreted as a variant of microfacet models with a particular microfacet distribution $p(\mathbf{h})$ and Fresnel term both explicitly present. Other terms ensure energy conservation and reasonable behavior of the model close to grazing angles. In “true” microfacet models, the shadowing/masking terms take on this responsibility.

The first and most important change which we make now is to allow an arbitrary normalized function $p(\mathbf{h})$ (with average \mathbf{h} vector in the direction of surface normal) instead of specific p_{Phong} . This change provides the needed flexibility and is in the spirit of [Ashikhmin et al. 2000] but preserves a much simpler mathematical formulation. Second, we replace C_0 continuous term $\max((\mathbf{k}_1\mathbf{n}), (\mathbf{k}_2\mathbf{n}))$ with $(\mathbf{k}_1\mathbf{n}) + (\mathbf{k}_2\mathbf{n}) - (\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})$. This modification is smoother and removes the banding artifacts of the A&S model which can appear in some situations. The specific shape of new term represents a bilinear patch with same values on the boundary of the domain as the original \max term (see Figure 3). Dropping the (\mathbf{kh}) term was found to produce a better overall appearance match to real-world

materials (see Section 3). We therefore arrive at

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{c p(\mathbf{h}) F((\mathbf{kh}))}{(\mathbf{k}_1\mathbf{n}) + (\mathbf{k}_2\mathbf{n}) - (\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})} \quad (2)$$

In agreement with our assertion that $p(\mathbf{h})$ has the largest impact, we found that the resulting appearance is not very sensitive to the exact form of the expression in the denominator. Similarly, instead of the full expression for the Fresnel term, we use Schlick’s simple polynomial approximation [1994]:

$$F((\mathbf{kh})) = r_0 + (1 - r_0)(1 - (\mathbf{kh}))^5 \quad (3)$$

where r_0 is reflectance at normal incidence. Since our model is not strictly physics-based, we prefer to think about this F term more as a handy empirical function controlling normal-to-grazing BRDF gain rather than assigning specific physical significance to it. This loose interpretation allows us, if necessary, to use r_0 values outside a rather narrow range that would be suggested by physical considerations (e.g., index of refraction) for a specific material.

2.1 General properties of the model

For any non-negative distribution $p(\mathbf{h})$ the model is clearly non-negative and reciprocal. The distribution directly affects the shape of the highlight allowing easy and intuitive modeling of its shape. Note that this is at most a 2D function (1D for isotropic surfaces) compared to the cumbersome four (three for isotropic) dimensions of full BRDFs. Although there is no distribution which exactly corresponds to a Lambertian BRDF, setting $p(\mathbf{h}) = \text{const}$ (and $F = 1$) results in surfaces very close to Lambertian in appearance making it possible to handle diffuse reflection within the same framework in most cases. The energy conservation condition states that for any \mathbf{k}_1 we should have

$$\int \rho(\mathbf{k}_1, \mathbf{k}_2) (\mathbf{k}_2\mathbf{n}) d\omega_{\mathbf{k}_2} \leq 1 \quad (4)$$

We substitute ρ from equation 2 and transfer the integration into the domain of half-vectors noting that [Torrance and Sparrow 1967] $d\omega_{\mathbf{k}_2} = 4(\mathbf{k}_1\mathbf{h})d\omega_{\mathbf{h}}$. Extending the new integral over the complete hemisphere of \mathbf{h} directions and noting that the denominator of equation 2 is not less than $(\mathbf{k}_2\mathbf{n})$ and $F \leq 1$, we arrive at a sufficient condition

$$4c \int p(\mathbf{h})(\mathbf{k}_1\mathbf{h})d\omega_{\mathbf{h}} \leq 1 \quad (5)$$

Since p is normalized, a universal but very loose constraint to ensure energy conservation is $c \leq 1/4$. A much tighter, but still conservative, bound for a specific distribution can be obtained by evaluating the maximal value of the integral in 5 over all \mathbf{k}_1 . For a typical BRDF with a single specular highlight near the mirror direction this maximal value is reached at $\mathbf{k}_1 = \mathbf{n}$, so in practice, only one integral needs to be computed.

In the following sections, we demonstrate the behavior of the proposed model in more detail and discuss its suitability for different applications.

3 Handling Measured Data

The most obvious way to use d-BRDF framework is to create new well-behaved BRDFs from user-specified distributions. However, if sufficiently high quality measurement data for a given material is available, a simple procedure exists to extract the necessary distribution from such data. In our experiments we used samples of




Figure 4: Backscattering data for metallic blue material. Different color channels are very close to being scaled versions of one another. Cross-section through the distribution is plotted for clarity (data for negative angles obtained by symmetry). Unnormalized Phong distributions with $n = 50$ and $n = 20$ are also shown.

different isotropic materials from Matusik et al. [2003a], Guenther et al. [2005] datasets and two samples from commercial service Integra. This selection covers a broad range of materials from highly specular to very diffuse and contains several paints, metals, plastic and fabrics. The dataset was recently augmented with data from a few anisotropic materials [Ngan et al. 2005].

To extract $p(\mathbf{h})$ from measured data, we note that for backscattering geometry $\mathbf{k}_1 = \mathbf{k}_2 \equiv \mathbf{k}$ we have $\mathbf{h} = \mathbf{k}$ and d-BRDF takes the form

$$\rho_{bs}(\mathbf{k}) \equiv \rho(\mathbf{k}, \mathbf{k}) = \frac{cr_0 p(\mathbf{k})}{2(\mathbf{k}\mathbf{n}) - (\mathbf{k}\mathbf{n})^2} \quad (6)$$

We can see that after multiplying $\rho_{bs}(\mathbf{k})$ by $2(\mathbf{k}\mathbf{n}) - (\mathbf{k}\mathbf{n})^2$ we can obtain a function proportional to the distribution $p(\mathbf{k})$. Note that we actually get three separate functions, one for each color channel. If a specific material is well represented by a single d-BRDF, these three functions should have the same shape but can differ by a constant factor, i.e. they should be scaled versions of one another (see Figure 4). One can then simply compute ratios of color components (e.g., green/red and green/blue) and average the three functions using these ratios as weights. After normalization, we obtain the distribution $p(\mathbf{h})$. The Fresnel parameter r_0 cannot be determined from backscattering data alone. One could perform a numerical fit to the data to obtain the “best” value of r_0 but we found that a much simpler procedure results in small difference in material’s appearance from that obtained through numerical fitting. First, we note that if r_0 is available, the remaining BRDF scaling constant c is trivially determined allowing us to start rendering images. We first rendered a test image with $r_0 = 1$ for all color channels. For many materials with no visually significant Fresnel behavior this was an immediate match. For other cases, we found it sufficient to set the ratio of color components or r_0 to be the same as that of the backscattering data so that only a single overall scaling coefficient has to be obtained to determine r_0 . Therefore, we simply render a few images with different r_0 scale and choose the one giving the best visual match to the original data. In tune with our loose non-physical interpretation of F , r_0 value can be also be directly estimated from the additional gain in strength $1/r_0$ of the specular highlight at close-to-grazing angles compared with normal incidence.

If backscattering functions for different color channels deviate sig-




Figure 6: Measured velvet BRDFs (left) and d-BRDFs constructed from $p(\mathbf{h}) = c_{norm}(1 + 4\exp(-c\tan(\theta)^2))$ distribution (right). Direct reconstruction of the distribution failed for this material.

nificantly from being just scaled versions of one another, diffuse component of ρ_{bs} should be subtracted out first. This corresponds to the standard diffuse plus specular separation. It is easy to do so by looking at the region outside the specular peak. If diffuse component is assumed to be Lambertian, the average value in this region can serve as that component of BRDF. For the non-Lambertian diffuse, a specific functional form should be chosen. For the single case we encountered which required such processing (plastic), the simple function $\rho_d(\mathbf{k}_1, \mathbf{k}_2) = d_0 + d_1(\mathbf{h}\mathbf{n})$ with two constants d_0 and d_1 gave a good match but we do not have sufficient data to recommend it as a general solution. Somewhat surprisingly, **for all other materials** the simplest procedure was sufficient, i.e., we did not have to separate out any special diffuse term, not even a Lambertian one.

Results for eight isotropic materials are shown on Figure 5 where side-by-side comparisons of measured BRDF with corresponding d-BRDF are shown for a simple scene. One material from the available dataset, a velvet-like fabric, an attempt to extract the distribution from the measured backscatter data directly through the process described above failed to produce any reasonable result due to very poor quality of the data close to retroreflection for this sample. However, as shown on Figure 6, a good visual match to this material was still obtained with a d-BRDF using simple analytic “inverted Gaussian with background” distribution [Ashikhmin et al. 2000] $p(\mathbf{h}) = c_{norm}(1 + 4\exp(-c\tan(\theta)^2))$ where θ is the angle between \mathbf{h} and surface normal. While some differences can be observed, we believe that the overall agreement is very good given the variety of material properties and the simplicity of the model and fitting procedures. As with most models, differences are the largest close to grazing angles but one should remember that typically the data are also least reliable there. Figure 7 shows a numerical comparison of the model results with measured data for metallic blue paint material. We plot the data on a logarithmic scale, which is often used as an approximation to the nonlinear perceptual response. While we could easily compute an error for our approximation, we still believe, however, that visual comparisons of the type presented on Figure 5 carry much more practical information than any existing error metrics.

To demonstrate that our approach is not restricted to 1D distributions, we applied the same extraction procedure to anisotropic materials. The main difficulty with available data [Ngan et al. 2005] in this case is its extreme sparsity, especially in the region of interest. To obtain a workable distribution, we therefore collected and averaged all valid samples within a 20 degree cone from exact retroreflection and further smoothed the distribution by applying symmetry considerations. As suggested by researchers who created the dataset, all points beyond 80 degree incident angle were considered unreliable. This region was linearly extrapolated (along θ angle) from neighboring data. The non-trivial distribution extracted from purple satin sample is shown on Figure 8, left. The image on the right was rendered with corresponding d-BRDF and demonstrates the characteristic complex shape of the highlight.




Figure 5: Measured BRDFs (left) and corresponding d-BRDFs (right). First column, from the top: aluminum bronze; metallic gold; silver paint; plastic. Second column, from the top: fabric2, synthetic fabric, cacao, rubber. No attempt was made to smooth either measured data or extracted distributions. For all materials except no explicit diffuse term was used.




Figure 7: Measured data for metallic blue material compared with d-BRDF values for different incident angles. Distribution was extracted directly from backscattering data. Top: scattering in the plane of incidence (θ_2 -dependence for $\phi_2 = \pi$). Bottom: ϕ_2 -dependence for constant θ_2 in the mirror direction ($\theta_2 = \theta_1$). Blue channel is shown on all plots.




Figure 8: Left: 2D distribution extracted from purple satin BRDF data. Right: resulting d-BRDF applied to a model with vertical anisotropy axis.

3.1 Implications for future measurement systems

To represent a material's full BRDF in our framework, we need only the backscattering data and some estimation of Fresnel parameter r_0 . This represents tremendous reduction in the amount of data which needs to be collected if the d-BRDF representation is acceptable for a given material. We do acknowledge, however, that measuring in this part of the domain is difficult with existing systems. The Matusik et al. [2003a] dataset is the only existing one in computer graphics containing sufficiently reliable backscattering data. Even for this excellent dataset, it was clear upon close examination that a significant part of the inaccuracies one can see on Figure 5 is due to problems with the data rather than inadequacy of d-BRDF representation. One common problem we encountered for highly specular materials is that due to geometric limitations on how close the light source and the detector can be positioned in backscattering configuration, the measurement system misses the most intense part of the specular highlight and therefore records a broader and less intense one. In a non-backscattering geometry, the detector can be positioned exactly in mirror reflection direction eliminating this problem. Data for anisotropic materials and a velvet-like fabric dis-




Figure 9: A slice through the center of a photograph (left) provides sufficient information to extract a distribution (right) which can be used to create a d-BRDF approximation, see the last image on Figure 10.

cussed above demonstrate more extreme cases of data-related problems when direct extraction of the distribution becomes impossible and either a much broader region of the domain or an appropriate analytic function have to be used.

To remove some of these problems, we believe that it is worth considering a special system for measuring backscattering. Since the light source and the detector for this geometry have to be close (in direction) and do not have to move with respect to each other, we believe that the resulting device is inherently more compact (and cheaper) than present systems. In fact, a common consumer camera with a built-in flash can serve as such a device. Although there is some separation between the lens and the flash, the fact that we were able to directly use Matusik’s data, which were measured with at least 5 degrees separation between the source and the detector, demonstrates that exact backscattering configuration is not crucial. A *single*, possibly high dynamic range, picture of known curved geometry (a cylinder is sufficient for isotropic materials) then provides enough data to extract the distribution. r_0 can then be estimated from the gain in highlight strength near grazing angle or from general considerations regarding particular material. We are not aware of any simple (i.e. not involving special-purpose measurement devices or radiometric calibration with known materials or light sources) technique to obtain *absolute* values for quantities related to light energy (luminance, irradiance, etc.) which are necessary to obtain the correct scaling factor c . However, in most applications it is common to set the overall scaling factor by hand anyway.

As a proof-of-concept experiment, we extracted a distribution from a low dynamic range photograph of a Christmas decoration shown on Figure 9. A single slice through the center of the ball was used. No correction for flash position or data smoothing were applied (except for averaging together left and right halves of the slice to obtain a symmetric distribution). The distribution was then used to render the image shown last on Figure 10 which presents satisfactory visual match to the real material. We believe that high dynamic range photography [Debevec and Malik 1997] combined with better raw data handling would allow higher quality measurements even with this exceptionally simple technique.

4 Hardware Implementation

The simple mathematical form of our model makes its implementation using modern graphics hardware straightforward with a technique similar to that used for factored representations [Kautz and McCool 1999; McCool et al. 2001]. Specifically, for each material, we create two textures: a two dimensional one to represent $p(\mathbf{h})$ and another one dimensional texture to hold Fresnel fraction $F((\mathbf{h}\mathbf{h}))$. In addition, we use a universal auxiliary texture holding

the term $(\mathbf{k}_1\mathbf{n})/((\mathbf{k}_1\mathbf{n}) + (\mathbf{k}_2\mathbf{n}) - (\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n}))$. In our Cg-based implementation, a vertex shader computes all needed texture coordinates given surface normal, light, and view directions. A fragment shader then performs texture lookups into the three textures, multiplies the results together and scales the product by color constant c and light intensity. On modern architectures multiple texturing units are available and the complete computation can be done in a single pass. Alternatively, simple analytic distributions as well as all other terms can also be evaluated directly inside the shader. The results are shown on Figure 10 and include both analytical distributions and those extracted from measured data as described in the previous section. Analytic examples were chosen to demonstrate versatility of the model and include an anisotropic distribution. They also show that a d-BRDF can be easily combined with a separate diffuse (in this case Lambertian) term.

Note that while values stored in Fresnel and auxiliary textures are guaranteed to lie in the $[0; 1]$ interval, this is not the case for distribution $p(\mathbf{h})$. This problem is not unique to our representation and always occurs when high dynamic range data are to be used in hardware. Different solutions have been proposed which range from simple rescaling and/or clamping (which we use) to more sophisticated techniques [Cohen et al. 2001]. Using a floating point framebuffer combined with a good tone mapping algorithm provides the best, although more complex, solution.

5 Efficient BRDF Sampling

Many modern rendering systems compute radiance leaving the surface through Monte-Carlo integration of the rendering equation [Kajiya 1986] $L(\mathbf{k}_2) = \int L(\mathbf{k}_1)p(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1\mathbf{n})d\omega_{\mathbf{k}_1}$. The importance sampling technique suggests that one should sample with probability density function (pdf) close in shape to the expression under the integral. However, the incident radiance is typically not known in advance, so sampling according to $p(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1\mathbf{n})$ is usually the best one can hope to do. One important desirable property of a reflection model is therefore the ability to generate samples with density proportional to (or at least sufficiently close to being proportional to) $p(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1\mathbf{n})$ for any given \mathbf{k}_2 .

For d-BRDFs we suggest a simple strategy of generating halfvectors \mathbf{h} with density $\text{pdf}(\mathbf{h}) = p(\mathbf{h})$ and then reflect the known vector \mathbf{k}_2 around \mathbf{h} to obtain sample direction \mathbf{k}_1 . If \mathbf{k}_1 is below horizon, we set the weight of such a sample to zero. Otherwise, the incoming luminance in the direction of the generated sample is weighted by the standard importance sampling expression

$$\frac{p(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1\mathbf{n})}{\text{pdf}(\mathbf{k}_2)} = \frac{4cF((\mathbf{h}\mathbf{h}))(\mathbf{k}_1\mathbf{n})(\mathbf{h}\mathbf{k}_1)}{(\mathbf{k}_1\mathbf{n}) + (\mathbf{k}_2\mathbf{n}) - (\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})} \quad (7)$$

where we used relationship $\text{pdf}(\mathbf{k}_2) = \text{pdf}(\mathbf{h})/4(\mathbf{k}_1\mathbf{h})$ between pdfs in \mathbf{h} - and \mathbf{k}_2 - spaces. To generate samples according to a general numerically represented $p(\mathbf{h})$, we use the standard cumulative density function (cdf) inversion procedure. For an isotropic distribution, we build a 1D cdf table and find \mathbf{h} ’s angle θ from the normal by solving $\xi = \text{cdf}(\theta)$ where ξ is a random variable uniformly distributed on $[0, 1]$. The azimuthal angle ϕ is sampled uniformly on $[0, 2\pi]$. Generalization to 2D distributions is straightforward and involves generating separate cdfs for each of a number of θ intervals. To generate an \mathbf{h} sample, one of these intervals is chosen first according to the distribution of total probability concentrated in it. Finally, ϕ is obtained using the corresponding interval-specific cdf. Further factorization into a (sum of) product of 1D factors can be attempted for 2D distributions, but we do not believe that the extra complexity is justified. If material’s BRDF is represented by a sum of several terms (for example, two d-BRDFs or a d-BRDF




Figure 10: Various BRDFs generated with our model, all implemented in hardware. Analytic distributions, top row: uniform; Phong with extra Lambertian term; synthetic velvet; anisotropic Phong ($n_1 = 1, n_2 = 100$) with two different anisotropy axis orientation. Approximations for measured materials, bottom row (see Figures 5 and 9): metallic blue; nickel; pearl paint; resin; Christmas ball.

and a Lambertian), the standard strategy is first to choose which term to sample and then to generate a sample according to the corresponding distribution. To make the choice, we can precompute directional hemispherical reflectance $R_i(\mathbf{k}_2)$ for each term i , i.e., the fraction of energy incident from a given direction \mathbf{k}_2 which is being reflected. A term is then chosen with probability proportional to its $R_i(\mathbf{k}_2)$. Note that overall pdf also becomes a weighted sum of pdfs of individual terms.

While expression 7 might not suggest this immediately (since ideal weight for each luminance sample should be 1), the proposed strategy is actually very efficient. This can be understood if we recall that the distribution has by far the greatest effect on most aspects of d-BRDF behavior. In fact, for measured data, not only one can sample the d-BRDFs approximation in this manner, but one can sample the *original* BRDF using the same sample generation procedure but different luminance sample weight

$$\frac{\rho_{\text{measured}}(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1 \mathbf{n})}{\text{pdf}(\mathbf{k}_2)} = \frac{4\rho_{\text{measured}}(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1 \mathbf{n})(\mathbf{k} \mathbf{h})}{p(\mathbf{h})} \quad (8)$$

where ρ_{measured} is the original BRDF. Complex analytical models (such as Cook-Torrance) can be sampled in a similar manner once an appropriate distribution is extracted from backscattering data¹. The result for metallic blue material is shown in the left column of Figure 11. Lawrence et al. [2004] recently proposed an efficient technique for sampling arbitrary BRDFs based on a special-purpose factorization procedure. The second column of Figure 11 shows the results obtained with this technique using the data available online for this material. Because of the different tone mapping procedures used, for all images on Figures ?? and 11 we multiplied the material BRDF by a factor of 10 to facilitate direct comparison with Figure 8 of [Lawrence et al. 2004]. Figure 12 shows a more complex scene with global illumination which uses many materials including measured data and the Cook-Torrance model. For both examples, we obtain sampling efficiency similar to that of the factored representation at a small fraction of the complexity. Interestingly,

¹This is the most general strategy effectively treating an analytic model as a measured one. For some models, a better distribution can probably be obtained using advanced knowledge about model details.

while the plastic material required special processing (separation of a non-Lambertian diffuse term) if a d-BRDF is to *replace measured data*, to obtain a distribution useful for *sampling original BRDF*, it was sufficient to follow the most basic extraction procedure with no separate diffuse term at all.

6 Discussion

The key advantage of the d-BRDF over most other simple models is that it does not restrict the distribution to any specific shape such as Phong or Gaussian. In many cases, the deviation of real distributions from these standard functions is quite significant. For example, the distribution on Figure 4 has a sharper central part and broader tails than either Phong or Gaussian. For many scenes, this discrepancy directly translates into noticeable visual differences and, in Monte Carlo systems, into reduced sampling efficiency.

Figure 13 compares d-BRDF representation with A&S model. The first row of images shows the importance of correct highlight shape. Phong distribution of A&S model has either a much sharper drop-off ($n = 50$) or significantly broader central part ($n = 20$) failing to match the details of appearance throughout the complete BRDF domain. Note that a separate diffuse component had to be used for A&S model - otherwise all regions outside the specular highlight would appear completely black. On the contrary, a single term d-BRDF was sufficient for this material. Realistic lighting environments and complex geometry used for images in the bottom row make it harder to identify specific problems due to material mismatch (most visible on the pedestal). Interestingly, while a d-BRDF provides a good much for both scenes, an $n = 50$ A&S material seems to us more acceptable for the sphere scene but $n = 20$ is a better match under more realistic conditions. Such context-sensitivity makes it very difficult to perform a reliable purely numerical fit using models which lack the necessary flexibility. The two rightmost columns of Figure 11 show the results of sampling original BRDF according to Phong distribution $p(\mathbf{h}) = (n+1)/(2\pi)\cos^n(\theta)$ with $n = 50$ and $n = 20$, which for this isotropic material this is the same as sampling according to A&S model. Either parameter choice re-




Figure 11: Sampling measured BRDF of metallic blue material with 16 (top row) and 100 (middle row) samples per pixel. From left to right: sampling according to extracted $p(\mathbf{h})$; sampling using Lawrence et al. factored representation; according to “best fit” Phong distribution ($n = 50$); according to Phong distribution with $n = 20$. Because of different tone mapping procedures, overall intensity was adjusted to better match Figure 8 of Lawrence et al.




Figure 12: A scene with global illumination rendered with our technique compared with factored representation of Lawrence et al. Materials of the table and the vase are represented by Cook-Torrance model while the bowl and the teapot use measured data (metallic blue, plastic, and nickel). 16 and 100 samples/pixel were used. Paths up to five bounces long were included.




Figure 13: *Top: single light scene. Bottom: complex model illuminated by a high dynamic range environment map (direct lighting only). Left to right: measured BRDF; d-BRDF using extracted distribution; A&S model with $n=50$ and $n=20$.*

sults in significant increase in noise levels compared with d-BRDF sampling.

Compared with previous approaches to handling measured data, we do not use any complex or potentially unstable numerical fitting procedures, and $p(\mathbf{h})$ is read almost directly from the data. The resulting representation is very compact: we store just 250 samples for the extracted isotropic distributions and 45×90 (extended to 45×180 by symmetry) samples for anisotropic ones. Furthermore, if necessary, the extracted distribution can be smoothed (and further compressed) by any standard techniques, for example by fitting it with polynomial splines for isotropic materials or spherical harmonics for anisotropic ones. This is much easier to do for the lower dimensional function $p(\mathbf{h})$ than for the original BRDF. Note that we do not have to use backscattering data to extract $p(\mathbf{h})$. Instead, data for any other incident angle can provide information about some part of $p(\mathbf{h})$ after corresponding processing. Combining such data for several incident angles will give a complete $p(\mathbf{h})$. One advantage of using backscattering data is that the distribution is read completely and most directly from it and distribution shape is not distorted by the Fresnel term with unknown r_0 . Another consequence of backscattering-based approach is that a very simple reflection measurement system becomes possible.

6.1 Limitations

One should not expect that the reflection properties of *all* materials can be accurately described by our model, or any simple model. The fact that we use a only a single low-dimensional function,

specifically that of the halfvector, although very convenient, prevents some effects from being accurately modeled. For example, materials which are designed to act as strong retroreflectors for a range of incident angles (such as those used in highway signs) can not be adequately represented. Effects due to wave properties of light are also not accounted for. We do believe though that d-BRDFs are adequate for a significant percentage of real-world materials. More general discussion on the subject of BRDF dimensionality reduction is given by Stark et al. [2005] and, from a different perspective, by Matusik et al. [2003a].

One can certainly create scenes where the relatively small differences between a real material and its d-BRDF become highly noticeable. In particular, since most of differences are concentrated near grazing angles, a scene with an infinite plane under low hanging light, as used by Stark et al. [2005], presents such case. If such configurations are of particular importance for a given application, the d-BRDF presents a unique ability to tailor the distribution to get closer to the needed appearance for specific conditions, most likely at the cost of worsening the agreement in other parts of the domain. As already mentioned, more realistic lighting condition or more complex geometry can make a precise BRDF match less crucial (see Figure 13 or compare images in Figure 11).

Finally, we caution against too direct a physical interpretation of the proposed model. While its general ideas are traceable to microfacet theory, the details lack a solid theoretical foundation leading to some conceptual problems if a physical interpretation is carried too far. In particular, the best values for the Fresnel parameter r_0 , while always inside the theoretically possible interval $[0, 1]$, often lie outside the rather narrow physically reasonable range for a given ma-

terial. We therefore suggest to think about F as an empirical term controlling normal-to-grazing BRDF gain which simply happens to behave like Fresnel. Keeping the (\mathbf{kh}) term in the denominator of equation 2 helps to make r_0 values somewhat more sensible. This form of the model can be helpful to increase reflection near grazing angle, but leads to worse color reproduction in that region. All examples shown do not use the (\mathbf{kh}) term.

7 Conclusion

We presented a simple BRDF model for visual computer graphics applications based on ideas from microfacet theory. The proposed representation is very flexible and possesses many desirable general properties. It is capable of providing a good approximation for measured reflection data with a simple fitting procedure, can be easily implemented in graphics hardware, and allows efficient sampling in a Monte-Carlo rendering system. Since the model uses only backscattering data to reconstruct the complete BRDF, it also suggests a new economical way of measuring surface reflection.

There are several possibilities for future research related to our model. First, our measurement experiment described in section 3.1 should not be considered as more than a proof of concept. We would like to build an actual system to measure backscattering with sufficient accuracy. We believe that this would make the collection of reflection data quicker and easier than is possible today. Once sufficient data is available, it would be interesting to perform an analysis of distributions similar to that done for full BRDFs by Matusik et al. [2003a] to isolate salient features responsible for specific appearance trends. Some common material properties can be incorporated by adding special terms to the model. For example, one designed for backscattering should be useful. Some systematic discrepancies between the proposed model and measured data near grazing angles suggest that a better overall model might be possible. Finally, more research is needed to establish robust practical BRDF difference metrics.

References

- ASHIKHMIN, M., AND SHIRLEY, P. S. 2000. An anisotropic phong brdf model. *Journal of Graphics Tools* 5, 2, 25–32.
- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based brdf generator. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 65–74.
- ASHIKHMIN, M., MARSCHNER, S., SHIRLEY, P., AND STAM, J., 2001. State of the art in measuring and modeling surface reflection. Siggraph Course Notes.
- BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. vol. 11, 192–198.
- COHEN, J., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2001. Real-time high-dynamic range texture mapping. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, 313–320.
- COOK, R. L., AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of SIGGRAPH 81)*, vol. 15, 307–316.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1 (Jan.), 1–34.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, 369–378.
- EDWARDS, D., BOULOS, S., JOHNSON, J., SHIRLEY, P., ASHIKHMIN, M., STARK, M., AND WYMAN, C. 2006. The halfway vector disk for brdf modeling. *ACM Transactions on Graphics* 25, 1 (Jan.), 1–18.
- FLEMING, R., DROR, R., AND ADELSON, E. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3, 5.
- GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P. S., ARVO, J. R., FERWERDA, J. A., PATTANAIK, S., LAFORTUNE, E. P. F., WALTER, B., FOO, S.-C., AND TRUMBORE, B. 1997. A framework for realistic image synthesis. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, 477–494.
- GÜNTHER, J., CHEN, T., GOESELE, M., WALD, I., AND SEIDEL, H.-P. 2005. Efficient acquisition and realistic rendering of car paint. In *Proceedings of 10th International Fall Workshop - Vision, Modeling, and Visualization (VMV) 2005*, Akademische Verlagsgesellschaft Aka GmbH, G. Greiner, J. Hornegger, H. Niemann, and M. Stamminger, Eds., 487–494.
- HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, vol. 25, 175–186.
- INTEGRA. www.integra.co.jp.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proceedings of SIGGRAPH 86)*, vol. 20, 143–150.
- KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary brdfs using separable approximations. In *Eurographics Rendering Workshop 1999*.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, 117–126.
- LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient brdf importance sampling using a factored representation. *ACM Transactions on Graphics* 23, 3 (Aug.), 496–505.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based brdf measurement including human skin. In *Eurographics Rendering Workshop 1999*.
- MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Transactions on Graphics* 24, 3 (Aug.), 727–734.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July), 759–769.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. Efficient isotropic brdf measurement. In *Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering*, 241–248.
- MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of brdfs for high-performance rendering. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 171–178.
- NEUMANN, L., NEUMANN, A., AND SZIRMAY-KALOS, L. 1999. Reflectance models with fast importance sampling. *Computer Graphics Forum* 18, 4 (Dec.), 249–265.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, 117–126.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I., AND LIMPERIS, T. 1977. Geometrical considerations and nomenclature for reflectance. Tech. Rep. 160, National Bureau of Standards.
- OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert's reflectance model. In *Proceedings of SIGGRAPH 94*, Computer Graphics Proceedings, Annual Conference Series, 239–246.
- PHONG, B.-T. 1975. Illumination for computer generated images. *Communications of the ACM* 18, 6, 311–317.
- SCHLICK, C. 1994. An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum* 13, 3, 233–246.
- STAM, J. 1999. Diffraction shaders. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, 101–110.
- STARK, M. M., ARVO, J., AND SMITS, B. 2005. Barycentric parameterizations for isotropic brdfs. *IEEE Transactions on Visualization and Computer Graphics* 11, 2 (March-April), 126–138.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surface. *Journal of Optical Society of America* 57, 9, 1105–1114.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, vol. 26, 265–272.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, vol. 26, 255–264.