Properties of Tabulated Bidirectional Reflectance Distribution Functions

Joel DeYoung and Alain Fournier
Department of Computer Science
University of British Columbia
201-2366 Main Mall
Vancouver, British Columbia V6T 1Z4 Canada
{deyoung,fournier}@cs.ubc.ca

Abstract

One way to overcome the limitations imposed by analytical models of reflection is to use discretely sampled reflectance data directly. Through either empirical measurement or simulation, a bidirectional reflectance distribution function (BRDF) is sampled and stored as a table of numbers. The generality of these tabulated BRDFs is useful for generating realistic images, but the inevitable inaccuracy associated with the data gathering process can lead to a BRDF that is much more general than it needs to be, or that lacks certain physical properties.

In this paper we propose measures for several properties of tabulated BRDFs: reciprocity, energy conservation, isotropy, and separability. Techniques to transform tabulated BRDFs to match one or more of these properties are also described. These transformations allow compression of the BRDF data, elimination of noise, improved computation time in some rendering tasks, and improved compliance with physical laws.

Une façon de pallier les lacunes imposées par les modèles de la réflexion analytique est d'utiliser directement les données de réflexion mesurées discrètement. Par la mesure empirique ou la simulation, une fonction de distribution de réflexion bidirectionnelle (FDRB) est acquise par échantillonage, et mise en mémoire comme une table des nombres. La généralité de ces FDRBs tabulaires est utile pour générer des images réalistes, mais l'inexactitude inévitable associée avec la méthode d'aquisition des données peut produire une FDRB qui est beaucoup plus générale que nécessaire, ou qui manque certaines propriétés physiques.

Cet article propose des mesures pour quelques propriétés des FDRBs tabulaires: la réciprocité, la conservation d'énergie, l'isotropie, et la séparabilité. Des techniques pour transformer les FDRBs tabulaires pourqu'elles possèdent ces propriétés sont aussi décrites. Ces transformations permettent la réduction des données de la FDRB, l'élimination du bruit, moins de temps de calcul pour produire les images, et une meilleure conformité aux lois physiques.

Keywords: Local illumination, BRDF, tabulated BRDF, singular value decomposition

1 Introduction

1.1 Bidirectional Reflectance Distribution Functions

Consider the geometry in Figure 1. Light arriving at a differential surface area dA through a solid angle $d\omega_i$ from direction $\vec{\omega_i}$ is reflected in some other direction $\vec{\omega_r}$. The amount of reflected radiance L_r is proportional to the incident irradiance E.

$$dL_r \propto dE \tag{1}$$

The constant of proportionality in Equation 1 is called the bidirectional reflectance distribution function, or BRDF. It is expressed as

$$f_r(\vec{\omega_i} \to \vec{\omega_r}) = \frac{dL_r}{dE} = \frac{dL_r}{L_i \cos\theta_i d\omega_i} \tag{2}$$

The BRDF is a function in four variables (a polar and azimuth angle for each of the incident and reflected directions). It is often written $f_r(\phi_i, \theta_i, \phi_r, \theta_r)$ to express this dependency explicitly.

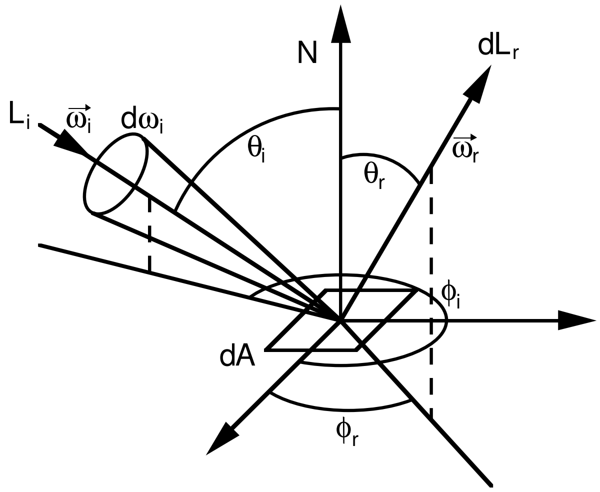

Figure 1: Shading Geometry

1.2 Tabulated BRDFs

In practice, the BRDF values are usually obtained from an analytical local illumination model. Most models in use are limited in the number of surfaces they can represent. There have been several recent attempts to obtain more realistic reflectance data through either empirical measurement or simulation. BRDF data obtained from such techniques can be fit to parameters of an illumination model[10], but can also be used directly[1, 4, 12]. In this paper we are concerned with the latter.

Since BRDFs of this type are usually in the form of a table of numbers, we term them *tabulated* BRDFs, or TBRDFs. Without loss of generality, we will assume in the rest of this paper that our data is monochromatic. In the case of irregular samples, each of the n_{f_r} samples will be characterized by a 5-tuple, the jth sample being:

$$(\phi_{i_j}, \theta_{i_j}, \phi_{r_j}, \theta_{r_j}, f_r(\phi_{i_j}, \theta_{i_j}, \phi_{r_j}, \theta_{r_j})) \qquad (3)$$

In the case of a regular sampling grid, the number of samples n_{f_r} in general is

$$n_{f_r} = n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r} \qquad (4)$$

where n_{ϕ_i}, n_{θ_i}, n_{ϕ_r}, and n_{θ_r} are the number of samples in the incident azimuth, incident polar, reflected azimuth, and reflected polar directions, respectively.

When obtaining tabulated BRDFs, measurement from physical samples is desirable, but is difficult since gonioreflectometers are expensive, quite slow, and the data from them is often quite noisy. Ward addressed the issues of cost and speed[10], but his technique results in nonuniform samples, complicating direct use of his data.

Cabral et al introduced the idea of sampling a BRDF by casting rays into a microgeometry[1]. This has become a popular technique for obtaining arbitrarily complex BRDFs[4, 12], since it is highly controllable, repeatable, and relatively fast. The BRDFs used as examples in this paper were either computed analytically or were obtained from such a virtual gonioreflectometer.

1.3 Motivation

Since TBRDFs are often obtained experimentally from real surfaces, they are much more general with regards to the kinds of surfaces they can represent. However, they suffer from the disadvantages associated with physical measurements, including both mechanically and experimentally induced errors. It would be useful to eliminate or reduce the impact of these errors before the measured BRDFs are used in rendering tasks.

For this reason, we are interested in examining properties of a BRDF that indicate how much error was introduced by the measurement process, so the BRDF data can be corrected. Since measured BRDFs are often very large, it would also be useful to examine properties of a BRDF that indicate how much accuracy would be lost if the BRDF data were compressed by making simplifying assumptions about its structure.

2 Measuring BRDF Properties

In this section, we present techniques for measuring four properties of tabulated BRDFs. The first two properties, reciprocity and energy conservation, express how close the TBRDF is to obeying physical laws. Lewis termed this *physical plausibility*[6]. Measured tabulated BRDFs *should* in theory be physically plausible by virtue of the way they are created. However, systematic errors and/or noise can be introduced by a gonioreflectometer, causing the resulting TBRDF to be implausible. TBRDFs obtained with a virtual gonioreflectometer may lack plausibility because of shortcomings in the underlying reflection model used to reflect light into the microgeometry. These measures, and the corresponding transformations presented in Section 3, are attempts to measure and eliminate the consequences of these phenomena.

The last two properties, isotropy and separability, can be thought of as measures of compressibility, since the associated transformations in Section 3 result in considerable space savings.

It should be noted that we are *not* attempting to measure how much effect these BRDF properties will have on generated images, but simply the properties themselves. This is an important distinction, since the former is so highly dependent on rendering parameters. The shape of an object, the lighting, and the viewing position can all greatly affect how much certain BRDF properties affect an image. Such an undertaking would necessitate a separate investigation.

For consistency, each of the measures has been formulated so that a value of zero implies the BRDF fully possesses the given property. Each of the measurements can vary from zero to infinity. The values obtained from these measurements are not dependent on the number of samples, meaning that measurements from two tabulated BRDFs of different resolutions are comparable.

2.1 Reciprocity

Helmholtz's reciprocity rule states that if the incident and reflected directions are reversed, the value of the BRDF should not change[9]. In the case where for every incident sample position (ϕ_i, θ_i) there is a reflected sample position (ϕ_r, θ_r) such that $\phi_i = \phi_r$ and $\theta_i = \theta_r$, then the reciprocity can be measured by averaging the square of the difference between all pairs of incident and reflected directions. This is achieved with the sum

$$P_r = \sqrt{\frac{\sum_{\phi_i} \sum_{\theta_i} \sum_{\phi_r} \sum_{\theta_r} \left[\begin{array}{c} f_r(\phi_i, \theta_i, \phi_r, \theta_r) - \\ f_r(\phi_r, \theta_r, \phi_i, \theta_i) \end{array} \right]^2}{2 n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{5}$$

Notice that we place an additional 2 in the denominator to account for the fact that each pair actually occurs twice in the sum, i.e. $f_r(\phi_i, \theta_i, \phi_r, \theta_r) - f_r(\phi_r, \theta_r, \phi_i, \theta_i)$ and $f_r(\phi_r, \theta_r, \phi_i, \theta_i) - f_r(\phi_i, \theta_i, \phi_r, \theta_r)$.

When the samples are on a regular grid, but the pairs are not necessarily matched, then to find the value $f_r(\phi_r, \theta_r, \phi_i, \theta_i)$ to match $f_r(\phi_i, \theta_i, \phi_r, \theta_r)$ we need to interpolate between samples. This can be most easily done using quadrilinear interpolation within the hypercube formed by the 16 samples surrounding the tuple $(\phi_r, \theta_r, \phi_i, \theta_i)$.

The irregular grid case leaves us with the difficult problem of interpolating between non-uniform samples. One technique is to take the j samples that are closest to the desired sample position (where j is sufficiently large) and to interpolate them with a suitable filter kernel. A useful choice in this case is a quadrivariate Gaussian with a standard deviation on the order of the mean distance of the sample points from the desired point.

It should be noted that an assumption of reciprocity is sometimes made when measuring BRDFs from physical samples in order to reduce the number of samples required. In this case, P_r will necessarily be zero.

2.2 Energy Conservation

The law of energy conservation says that the total amount of exitance M must be less than or equal to the total amount of incident irradiance E. For a given incident direction $\vec{\omega_i}$, the ratio of M to E is expressed as

$$\frac{M}{E} = \int_{\Omega_N} f_r(\vec{\omega_i} \rightarrow \vec{\omega_r}) \, (N \cdot \vec{\omega_r}) \, d\omega_r \tag{6}$$

For a BRDF f_r to conserve energy, $\frac{M}{E}$ must be less than or equal to one for all possible values of $\vec{\omega_i}$.

Notice that an integral over the hemisphere is equivalent to a double integral over the polar and azimuth angles. Recall that the integral in Equation 6 is in terms of the reflected solid angle $d\omega_r$, which is not uniform as θ_r varies. Noting that $d\omega_r = d\theta_r d\phi_r \sin\theta_r$ and that $(N \cdot \vec{\omega_r}) = \cos\theta_r$, we can express the integral as

$$\frac{M}{E} = \int_{\phi_r} \int_{\theta_r} f_r(\phi_i, \theta_i, \phi_r, \theta_r) \sin\theta_r \cos\theta_r \; d\theta_r \; d\phi_r \tag{7}$$

When the samples are uniform in the θ and ϕ directions, we can easily approximate the integral with a sum. For consistency, we subtract one from the resulting integral so that energy conserving behavior will be expressed as zero. This yields

$$\Gamma(\phi_i, \theta_i) = \max(0, \sum_{\phi_r} \sum_{\theta_r} \left[\begin{array}{c} f_r(\phi_i, \theta_i, \phi_r, \theta_r) \cdot \\ \sin\theta_r \cos\theta_r \cdot \\ \Delta\phi_r \Delta\theta_r \end{array} \right] - 1) \tag{8}$$

where $\Delta\phi_r$ and $\Delta\theta_r$ are the radian distance between samples in the reflected ϕ and θ directions respectively.

Since we are interested in energy conservation of the entire BRDF, we take the average of all Γ values, giving us our final measure.

$$P_{ec} = \frac{\sum_{\phi_i} \sum_{\theta_i} \Gamma(\phi_i, \theta_i)}{n_{\phi_i} n_{\theta_i}} \tag{9}$$

Notice that there is no need to assume that the samples are uniform in the θ_i and ϕ_i directions.

When the samples are irregular in all directions, we need another way to approximate the integral. A good method in this case is to triangulate in 2D the set of samples for a given incident direction $\vec{\omega_i}$. For each of these triangles, Equation 6 is the "differential-to-finite form factor" of the triangle with respect to the infinitesimal surface element dA. This can be computed using classic methods such as the Nusselt analog. The integral $\Gamma(\phi_i, \theta_i)$ is then obtained by summing the terms from all the triangles.

2.3 Isotropy

In general, BRDFs as described in Equation 2 are *anisotropic*, meaning that the reflected radiance varies as the surface is rotated about the normal N. In contrast, an isotropic BRDF remains constant under such a rotation, and can be expressed as

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = f_r^{iso}(\theta_i, \phi_r - \phi_i, \theta_r) \qquad (10)$$

This simplification can reduce storage and computational costs considerably. In the case of regular samples, this reduces the total number of samples in Equation 4 to

$$n_{f_r} = n_{\theta_i} n_{\phi_r} n_{\theta_r} \qquad (11)$$

To measure the isotropy of a tabulated BRDF in the simple case where $n_{\phi_i} = n_{\phi_r}$, we start by computing the average BRDF value for the set of angles where the difference between ϕ_i and ϕ_r is the same. We use the shorthand $\phi_+ = \phi_r - \phi_i$ to denote this difference.

$$\mu_i(\theta_i, \phi_+, \theta_r) = \frac{\sum_{\phi_i} f_r(\phi_i, \theta_i, \phi_i + \phi_+, \theta_r)}{n_{\phi_i}} \qquad (12)$$

We are interested in the average deviation from the mean μ_i, expressed as

$$\sigma_i(\theta_i, \phi_+, \theta_r) = \sqrt{\frac{\sum_{\phi_i} \left[\begin{array}{c} f_r(\phi_i, \theta_i, \phi_i + \phi_+, \theta_r) - \\ \mu_i(\theta_i, \phi_+, \theta_r) \end{array} \right]^2}{n_{\phi_i}}} \qquad (13)$$

This is equivalent to taking the standard deviation of all BRDF values for a given $(\theta_i, \phi_+, \theta_r)$.

To obtain our final measure of isotropy, we simply average the σ_i values over all possible values of $(\theta_i, \phi_+, \theta_r)$

$$P_i = \frac{\sum_{\theta_i} \sum_{\phi_+} \sum_{\theta_r} \sigma_i(\theta_i, \phi_+, \theta_r)}{n_{\theta_i} n_{\phi_r} n_{\theta_r}} \qquad (14)$$

In the irregular grid case, it is necessary to choose a set of ϕ_+ values, and for each sample value compute $f_r(\phi_i, \theta_i, \phi_i + \phi_+, \theta_r)$ using the interpolation methods described in Section 2.1. The resulting sets of BRDF values can be used in the computation of μ_i in Equation 12.

2.4 Separability

To say that a BRDF is separable, we mean that it can be represented as a product of two functions, one for the incident light and one for the reflected light.

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = f_r^{in}(\phi_i, \theta_i) \times f_r^{out}(\phi_r, \theta_r) \qquad (15)$$

This removes the cross-dependence between the incident and reflected directions, thereby requiring much less storage space.

$$n_{f_r} = n_{\phi_i} n_{\theta_i} + n_{\phi_r} n_{\theta_r} \qquad (16)$$

An additional benefit is that radiosity computations, like those described by Cohen and Wallace[2], can be done efficiently with non-Lambertian surfaces[7].

For the discussion of separability, we will assume that for every pair (ϕ_i, θ_i) there is a pair (ϕ_r, θ_r) with the same direction. If this is not the case, such as in some experimental measurements, we have to generate the missing pairs using interpolation, as described in Section 2.1. Of course if we assume that reciprocity holds, then the required values can be generated trivially.

Consider a $Q \times Q$ matrix A where the rows correspond to a lexicographic ordering of the pair (ϕ_i, θ_i), and the columns correspond to a lexicographic ordering of the pair (ϕ_r, θ_r)[1]. If we then fill the matrix A with the BRDF values, we can think of the two separated functions f_r^{in} and f_r^{out} as two vectors u and v whose outer product results in A. To find u and v, we must factor A. The singular value decomposition (SVD) numerical technique is ideal for this, as it allows us to derive a measure of how close to factorable A is, as well as actually determine the most probable values for u and v (Section 3.4)[2].

[1] Note that in the case of a reciprocal TBRDF, the matrix A will be symmetric.

[2] For more details on the SVD, the reader is directed to the book by Watkins[11].

The singular value decomposition separates A into three matrices.

$$A = UDV^T \tag{17}$$

In the above equation, U, V^T and D are all $Q \times Q$ matrices. D is a diagonal matrix containing the singular values of A (in descending order). The number of non-zeros in D specifies the *rank* of A. For us, the rank is important because it indirectly indicates how close A is to being separable. If A has a rank of one, meaning that D is zero except for d_{11}, then there is an exact solution for u and v, and the BRDF is separable. To measure how close to separable the function is, we define another $Q \times Q$ matrix $D^{(1)}$.

$$D^{(1)} = \begin{bmatrix} d_{11} & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{18}$$

We then substitute $D^{(1)}$ for D and re-multiply to obtain $A^{(1)}$, the matrix of rank one based on A. Notice that the sparse nature of $D^{(1)}$ allows this multiplication to be performed quickly.

$$A^{(1)} = UD^{(1)}V^T \tag{19}$$

The magnitude of the difference between $A^{(1)}$ and A gives us a measure of separability. We measure this magnitude similar to the Frobenius norm of the difference matrix, except that we divide each square by the total number of elements in the matrix. This removes the dependence of P_s on the number of samples.

$$P_s = \sqrt{\sum_p \sum_q \frac{[A^{(1)}_{pq} - A_{pq}]^2}{n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{20}$$

In the SVD, A is expressed as a weighted sum of outer products of the columns of U and the rows of V^T, with the singular values acting as weights. Equation 20 only deals with expressing A as the product of the first such column and row. Fournier showed in [3] that expressing a BRDF as a sum of separable function products is desirable: the advantage of improved radiosity computation is not lost, and a larger subset of all BRDFs can be separated. If a BRDF can be expressed as a sum of k separated functions, we say that it is k-separable. The k-separated BRDF is reconstructed as:

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = \sum_{j=1}^{k} [f_r^{in(j)} \times f_r^{out(j)}] \tag{21}$$

To measure k-separability, we find the closest matrix of rank k to A. This is done by removing all but the k largest singular values from D.

$$D^{(k)} = \begin{bmatrix} d_{11} & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & d_{22} & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & \ddots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & d_{kk} & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ \vdots & & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \tag{22}$$

We then re-multiply to get the rank k matrix $A^{(k)}$.

$$A^{(k)} = UD^{(k)}V^T \tag{23}$$

This allows us to measure the k-separability $P_s^{(k)}$.

$$P_s^{(k)} = \sqrt{\sum_p \sum_q \frac{[A^{(k)}_{pq} - A_{pq}]^2}{n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{24}$$

3 BRDF Transformations

We may wish to change a tabulated BRDF to match one or more of the four properties covered in the previous section. Doing so with the properties of reciprocity and energy conservation makes the BRDF physically plausible. Doing so with the properties of isotropy and separability compresses the BRDF by changing the data to fit these two simplifying assumptions.

We may also wish to change the BRDF to be *closer* to a certain property, placing it intuitively in between the original BRDF and one matching the property. For this purpose, we define a constant δ, which indicates how close to a given property the BRDF should be. In each section, we will first show how to transform a BRDF to match the property ($\delta = 1$), and then generalize the technique to allow for more gradual transformations ($0 < \delta < 1$).

If the values of the TBRDF are obtained experimentally, then they may come with an error estimate. In this case all of the linear combinations used below should be changed to weighted sums, where the normalized weights are inversely proportional to the estimate of the errors at each data point.

3.1 Reciprocity

To make a tabulated BRDF reciprocal, we simply set each value to the average of itself and its reciprocal.

$$f_r^r(\phi_i,\theta_i,\phi_r,\theta_r) = \frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r) + f_r(\phi_r,\theta_r,\phi_i,\theta_i)}{2} \tag{25}$$

To make the BRDF *more* reciprocal, we linearly interpolate between a given sample and the average between that sample and its reciprocal. This gives us

$$\begin{aligned} f_r'(\phi_i,\theta_i,\phi_r,\theta_r,\delta) = \\ (1-\delta) f_r(\phi_i,\theta_i,\phi_r,\theta_r) + \\ \delta f_r^r(\phi_i,\theta_i,\phi_r,\theta_r) \end{aligned} \tag{26}$$

As mentioned above, these values can be weighted according to their error estimates. This is especially important in the case of irregular samples, since in general one value will be measured and the other will be the result of interpolation, giving it a different confidence interval. Estimating the confidence interval for the interpolated values is a non-trivial problem without an underlying reflection model, since a model would have to be available to determine the confidence interval at an arbitrary sample position.

3.2 Energy Conservation

If a tabulated BRDF does not conserve energy, it is because the total energy reflected is greater than the total energy received. It therefore seems natural to reduce every BRDF value associated with a given incident direction by a scalar sufficient to ensure energy conservation. For a given incident direction, this can be thought of as uniformly reducing the "size" of the BRDF.

Recall the definition of $\Gamma(\phi_i,\theta_i)$ in Equation 8. We wish to divide each BRDF value by a constant factor x that will make $\Gamma(\phi_i,\theta_i)$ be exactly zero.

$$\sum_{\phi_r}\sum_{\theta_r}\left[\frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{x}\sin\theta_r\cos\theta_r\ \Delta\phi_r\Delta\theta_r\right] - 1 = 0 \tag{27}$$

$$\frac{1}{x}\sum_{\phi_r}\sum_{\theta_r}[f_r(\phi_i,\theta_i,\phi_r,\theta_r)\sin\theta_r\cos\theta_r\ \Delta\phi_r\Delta\theta_r] = 1 \tag{28}$$

$$\frac{1}{x} = \frac{1}{\sum_{\phi_r}\sum_{\theta_r}[f_r(\phi_i,\theta_i,\phi_r,\theta_r)\sin\theta_r\cos\theta_r\ \Delta\phi_r\Delta\theta_r]} \tag{29}$$

$$x = \Gamma(\phi_i,\theta_i) + 1 \tag{30}$$

Therefore for a given incident direction, our factor is simply $\Gamma(\phi_i,\theta_i)+1$, resulting in the definition for an energy conserving BRDF.

$$f_r^{ec}(\phi_i,\theta_i,\phi_r,\theta_r) = \frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{\Gamma(\phi_i,\theta_i)+1} \tag{31}$$

To transform a tabulated BRDF to be *closer* to conserving energy, we simply multiply $\Gamma(\phi_i,\theta_i)$ by δ.

$$f_r'(\phi_i,\theta_i,\phi_r,\theta_r,\delta) = \frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{\delta\Gamma(\phi_i,\theta_i)+1} \tag{32}$$

It could be argued that this technique is somewhat artificial. It only changes those incident directions that violate energy conservation by clamping the data to the maximum that is theoretically plausible. There are two main problems with this. First, the scaling operations that are performed are not uniform since the scalar is different for each incident direction. Second, scaling the data to make $\frac{M}{E}$ have a value of exactly one may obey physical laws, but the case could be made that it is not realistic. No surfaces in the real world exhibit such ideal behavior. Even the most perfect reflectors ever observed still have values of $\frac{M}{E}$ less than one.[5].

The first concern can be addressed by choosing one scalar value for the entire BRDF. This scalar would be chosen based on the maximum of all values of $\Gamma(\phi_i,\theta_i)$, changing Equation 31 to

$$f_r^{ec}(\phi_i,\theta_i,\phi_r,\theta_r) = \frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{\max_{(\phi_i,\theta_i)}[\Gamma(\phi_i,\theta_i)]+1} \tag{33}$$

The second concern can be dealt with by clamping $\frac{M}{E}$ to some value τ such that $0 < \tau <= 1$. This changes the definition of $\Gamma(\phi_i,\theta_i)$ in Equation 8 to

$$\Gamma(\phi_i,\theta_i) = \max(0, \sum_{\phi_r}\sum_{\theta_r}\left[\begin{array}{c} f_r(\phi_i,\theta_i,\phi_r,\theta_r)\cdot \\ \sin\theta_r\cos\theta_r\cdot \\ \Delta\phi_r\Delta\theta_r \end{array}\right] - \tau) \tag{34}$$

The energy conserving BRDF is then computed by changing Equation 31 to

$$f_r^{ec}(\phi_i,\theta_i,\phi_r,\theta_r) = \frac{\tau\cdot f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{\Gamma(\phi_i,\theta_i)+\tau} \tag{35}$$

3.3 Isotropy

Transforming a tabulated BRDF to be isotropic is very simple. The average BRDF value μ_i from Equation 12 is used for all values with the same difference between ϕ_r and ϕ_i. This yields

$$f_r^{iso}(\theta_i,\phi_+,\theta_r) = \mu_i(\theta_i,\phi_+,\theta_r) \tag{36}$$

To make a tabulated BRDF *more* isotropic, the value is linearly interpolated between f_r and μ_i. This results in

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = (1-\delta) f_r(\phi_i, \theta_i, \phi_r, \theta_r) + \delta \mu_i(\theta_i, \phi_r - \phi_i, \theta_r) \quad (37)$$

3.4 Separability

Recall from Section 2.4 that separating the BRDF is equivalent to factoring a matrix A into the outer product of two vectors u and v. When all but the largest of the singular values in D are zeroed, it becomes evident that the only numbers affecting $A^{(1)}$ are the first column of U and the first row of V^T. Therefore, u is taken to be the first column of U, and v the first row of V^T. This gives a definition for f_r^{in} and f_r^{out}[3].

$$f_r^{in}(\phi_i, \theta_i) = d_{11} u_p \quad (38)$$

$$f_r^{out}(\phi_r, \theta_r) = v_q \quad (39)$$

Finding the kth separable component simply involves using the kth column of U, $u^{(k)}$ and the kth row of V^T, $v^{(k)}$.

$$f_r^{in(k)}(\phi_i, \theta_i) = d_{kk} u_p^{(k)} \quad (40)$$

$$f_r^{out(k)}(\phi_r, \theta_r) = v_q^{(k)} \quad (41)$$

This is identical to Fournier's technique of separating a BRDF into a sum of k separable models[3].

If a value of $\delta < 1$ is chosen in order to make a BRDF closer to k-separable, it is obvious that f_r cannot be separated into f_r^{in} and f_r^{out}, but the partially transformed BRDF is expressed as

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = (1-\delta) A_{pq} + \delta A_{pq}^{(k)} \quad (42)$$

4 Experiments and Results

Figure 3 shows an example of performing a reciprocity transform on a Phong BRDF where $P_r = 0.117646$. Lewis showed in [6] that Phong shaders are never reciprocal. Performing the transformation resulted in a noticeable change in the BRDF data. However, differences in P_r do not seem to affect rendered images a great deal.

Figure 4 demonstrates transforming a BRDF to conserve energy. At the left is a plot of a Phong BRDF where $P_{ec} = 0.160301$ ($\theta_i = 50°$), and to the right of that is the energy conserving version. The reduction in size is more subtle than in the two rightmost plots, where the incident polar angle is larger ($\theta_i = 90°$). This is exactly what is expected since Phong's energy conserving behavior becomes worse at large incident polar angles. Again, the visual difference in rendered images is usually negligible. However, the underlying difference between the two BRDFs is very important, especially if the BRDFs are used in global illumination computations.

Figure 5 shows an example of transforming a BRDF to be isotropic. This anisotropic BRDF that simulates brushed metal was generated by running a microgeometry of parallel cylinders through a virtual gonioreflectometer[4]. The resulting isotropy measurement was $P_i = 0.0319745$. On the left teapot, notice the specular highlight running along the base. This corresponds to the scratches in the brushed metal surface. On the right teapot, the BRDF has been transformed to be completely isotropic ($\delta = 1.0$). Notice the difference in the specular highlight. Figure 6 shows plots of several scratched metal BRDFs. The one on the left is the original BRDF used in Figure 5. The one on the right is the isotropic version used in Figure 5. The middle one is halfway between the other two ($\delta = 0.5$).

Figure 7 demonstrates separating a BRDF. The microgeometry on the left was used to obtain a velvet BRDF, which was used to render the three chairs shown. The left chair was rendered with the original BRDF. The BRDF was then separated and used to render the middle chair. The right chair was rendered with a version made up of a sum of five separated components. Notice that this version is virtually indistinguishable from the original, yet the BRDF only requires 4.4% of the storage space! Figure 2 shows the value of $P_s^{(k)}$ for the velvet BRDF as k increases. Curves for other BRDFs exhibited very similar behavior.

5 Conclusion and Future Research

In this paper we have proposed techniques for measuring four properties of BRDFs: reciprocity, energy conservation, isotropy, and separability. As well, we have presented methods for transforming BRDF

[3] In Equations 38 through 42, $p = \phi_i n_{\theta_i} + \theta_i$ and $q = \phi_r n_{\theta_r} + \theta_r$.

[4] The use of parallel cylinders to simulate brushed metal was inspired by the work of Poulin and Fournier[8].

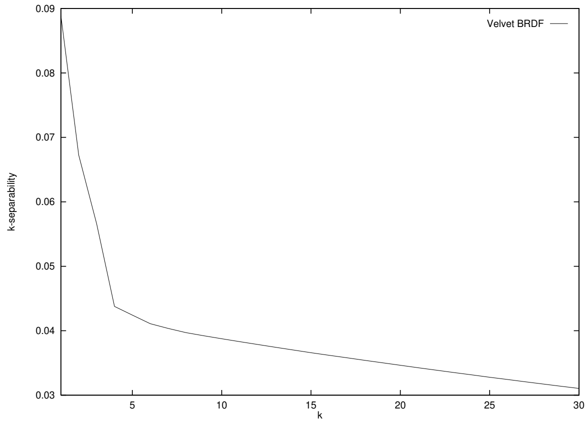

Figure 2: Measure of k-separability of the velvet BRDF for $1 \leq k \leq 30$.

data to possess these properties.

There are several useful features of these techniques. First, the BRDF data can be compressed by exploiting the properties of isotropy and separability. Second, noise introduced into the data by a gonioreflectometer, or errors caused by the shortcomings of a virtual gonioreflectometer, can be eliminated by ensuring that the BRDF be reciprocal and conserve energy. Third, physical plausibility of a BRDF can be guaranteed in situations where these properties are a prerequisite to rendering. Finally, transforming a BRDF to be separable allows radiosity computations to be performed efficiently for non-Lambertian surfaces, allowing more complex surfaces to be rendered using this popular technique.

More experimentation is necessary to fully investigate this approach, most notably on the type of data supplied by Ward[10].

References

[1] Brian Cabral, Nelson Max, and Rebecca Springmeyer. Bidirectional reflection functions from surface bump maps. In *Computer Graphics (SIGGRAPH '87 proceedings)*, pages 273–281, 1987.

[2] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis.* Academic Press Professional, 1993.

[3] Alain Fournier. Separating reflection functions for linear radiosity. In *Sixth EUROGRAPHICS Workshop on Rendering*, pages 383–392, 1995.

[4] Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. Wavelength dependent reflectance functions. In *Computer Graphics (SIGGRAPH '94 proceedings)*, pages 213–220, 1994.

[5] Eugene Hecht. *Optics.* Addison-Wesley Publishing Company, 1987.

[6] Robert R. Lewis. Making shaders more physically plausible. In *Fourth EUROGRAPHICS Workshop on Rendering*, pages 47–61, 1993.

[7] László Neumann and Attila Neumann. Photosimulation: Interreflection with arbitrary reflectance models and illumination. *Computer Graphics Forum*, 8:21–34, March 1989.

[8] Pierre Poulin and Alain Fournier. A model for anisotropic reflection. In *Computer Graphics (SIGGRAPH '90 proceedings)*, pages 273–281, 1990.

[9] Hermann von Helmholtz. *Helmholtz's Treatise on Physiological Optics.* Optical Society of America, 1924.

[10] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Computer Graphics (SIGGRAPH '92 proceedings)*, pages 265–272, 1992.

[11] David S. Watkins. *Fundamentals of Matrix Computations.* John Wiley and Sons, Inc., 1991.

[12] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. In *Computer Graphics (SIGGRAPH '92 proceedings)*, pages 255–264, 1992.

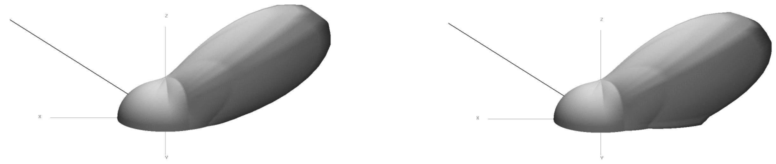

Figure 3: Two plots of a Phong BRDF ($\phi_i = 0°, \theta_i = 50°$), before and after a reciprocity transformation

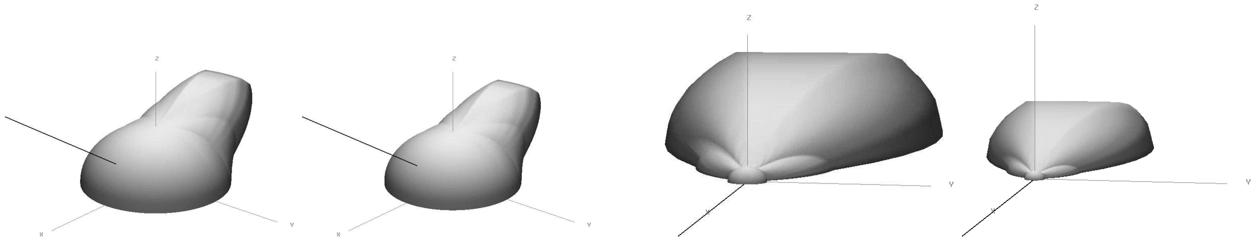

Figure 4: Four plots of a Phong BRDF. The leftmost two show the BRDF for the incident direction of $\phi_i = 0°, \theta_i = 50°$, before and after an energy conservation transformation. The rightmost two show the same BRDF for the incident direction of $\phi_i = 0°, \theta_i = 90°$, before and after the same transformation.

Figure 5: Two teapots rendered with an anisotropic BRDF representing brushed metal and the transformed isotropic version

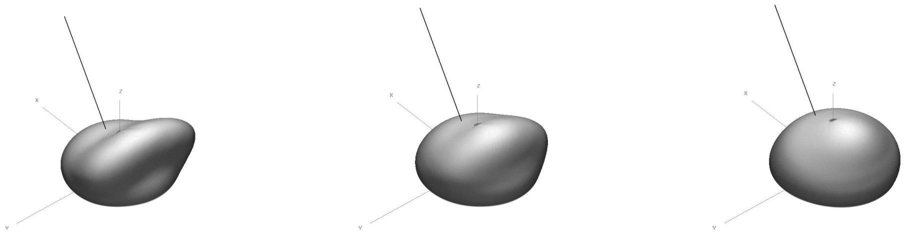

Figure 6: Plots of three brushed metal BRDFs ($\phi_i = 0°$, $\theta_i = 20°$) transformed to be isotropic with different δ values

Figure 7: Velvet microgeometry and three velvet chairs