Efficient BRDF Importance Sampling Using A Factored Representation

Jason Lawrence[*]
Princeton University

Szymon Rusinkiewicz[†]
Princeton University

Ravi Ramamoorthi[‡]
Columbia University

Abstract

High-quality Monte Carlo image synthesis requires the ability to importance sample realistic BRDF models. However, analytic sampling algorithms exist only for the Phong model and its derivatives such as Lafortune and Blinn-Phong. This paper demonstrates an importance sampling technique for a wide range of BRDFs, including complex analytic models such as Cook-Torrance and measured materials, which are being increasingly used for realistic image synthesis. Our approach is based on a compact factored representation of the BRDF that is optimized for sampling. We show that our algorithm consistently offers better efficiency than alternatives that involve fitting and sampling a Lafortune or Blinn-Phong lobe, and is more compact than sampling strategies based on tabulating the full BRDF. We are able to efficiently create images involving multiple measured and analytic BRDFs, under both complex direct lighting and global illumination.

Keywords: BRDF, Importance Sampling, Monte Carlo Integration, Global Illumination, Rendering, Ray Tracing

1 Introduction

The modeling of complex appearance is a key component in using photo-realistic rendering techniques to produce convincing images. The subtleties of how light interacts with different surfaces provide important clues about material and finish—impressions that cannot be conveyed by geometry alone. In this paper, we focus on one component of appearance models, the Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977]. A major challenge in incorporating complex BRDF models into a Monte Carlo-based global illumination system is efficiency in sampling: when tracing a path through the scene, it is desirable to generate reflected rays distributed according to the BRDF. When simulating light reflecting from a mirror-like surface, for example, most of the energy will be in rays close to the direction of ideal specular reflection. In this situation, it is wasteful to generate reflected rays in random directions: many fewer total paths will be necessary if the rays are generated mostly in the specular direction. This technique, known as importance sampling, reduces image variance and is critical for efficient rendering.

Effective importance sampling strategies are known only for the simplest Lambertian and Phong models, and generalizations such as Lafortune's cosine lobes [1997]. More complex BRDFs, including both measured data and physically-based analytic models (such as Cook-Torrance [1982], which has been used for over 20 years) have no corresponding importance sampling strategies. This difficulty has impeded the widespread adoption of realistic BRDFs in image synthesis.

[*]e-mail: jlawrenc@cs.princeton.edu

[†]e-mail: smr@cs.princeton.edu

[‡]e-mail: ravir@rendering.cs.columbia.edu

(a) Lafortune sampling   (b) Our method

Figure 1: *Monte Carlo renderings of a vase with a shiny Cook-Torrance BRDF under complex illumination. These images were generated with a path tracer that selected 16 reflected rays per pixel according to (a) a best-fit Lafortune model with 3 lobes, and (b) our factored BRDF representation. Notice the reduced noise in the image on the right that was made with our approach: the other algorithm would need roughly 6 times more samples to achieve the same quality.*

This paper develops a general, efficient importance sampling algorithm for arbitrary BRDFs. Our algorithm is based on *reparameterizing* the BRDF, followed by a compact and accurate *decomposition* into factors. We express the four-dimensional BRDF as a sum of (a small number of) terms, each of which is the product of a two-dimensional function that depends only on view direction, and two one-dimensional functions. This representation allows importance sampling by numerical inversion of the Cumulative Distribution Functions of the 1D factors. In contrast with other BRDF factorization techniques [Kautz and McCool 1999; McCool et al. 2001], which are geared towards real-time rendering, our representation is developed especially for sampling. To this end, we present the first factored BRDF representation that can be directly sampled. We demonstrate the benefits of our method with examples in which the variance of a Monte Carlo estimator is reduced by a factor of 2-10 over the best alternative sampling strategies for both analytic models, such as Cook-Torrance (see Figure 1), and measured materials, including the BRDFs acquired by Matusik et al. [2003].

2 Related Work

The goal of global illumination algorithms is numerical integration of the rendering equation, first formulated by Kajiya [1986]:

$$L_o(x, \omega_o) = L_e(x, \omega_o) + \int_{\Omega_{2\pi}} L_i(x, \omega_i)\, \rho(x, \omega_i, \omega_o)\, (\omega_i \cdot \mathbf{n})\, d\omega_i.$$

The approach of Monte Carlo algorithms is to evaluate the incoming radiance L_i by recursively casting rays through the scene to simulate light transport paths.

In order to reduce variance, it is desirable to *importance sample* reflected rays by preferentially considering paths carrying high energy. One possibility is to importance sample according to the lighting L_i, as in sampling light sources for direct illumination. In the special case when L_i is described by an environment map, effective illumination sampling methods have recently been demonstrated [Agarwal et al. 2003; Kollig and Keller 2003]. These methods are intended for diffuse or rough semi-glossy materials, and typically require a few hundred samples for good results. At low to medium sample counts, they may miss important details, especially for glossy materials and slowly-varying environments.

In the general case of global illumination (Figure 10), it is not practical to pre-compute the incoming radiance at all points in the scene. Therefore, we seek to sample according to the product of the BRDF ρ and the incident cosine term $\omega_i \cdot \mathbf{n}$. The remainder of this section describes related work on obtaining the complex analytic and measured BRDFs we would like to sample, deriving factored representations, and performing importance sampling in the context of global illumination.

BRDF models: There exist many analytical models for the BRDF that approximate the way specific materials reflect light. Some of these are phenomenological, such as the popular Phong shading model [Phong 1975]. More sophisticated, physically-based analytical models can capture effects including Fresnel reflection and rough microgeometry [Torrance and Sparrow 1967; Cook and Torrance 1982; He et al. 1991]. Anisotropic reflection models characterizing the reflective properties of oriented surfaces such as brushed metal have also been developed [Kajiya 1985; Ward 1992; Poulin and Fournier 1990; Ashikhmin and Shirley 2000]. Other analytical BRDF models, such as those meant to describe dusty surfaces, exhibit backscattering phenomena [Hapke 1963; Oren and Nayar 1994]. Despite the large amount of research on these BRDFs, most of these models have so far been difficult to sample efficiently. This typically arises because the analytic formula is difficult or impossible to integrate and invert.

The potential benefit of using measurements of a BRDF has also gained recent attention [Ward 1992; Greenberg et al. 1997; Dana et al. 1999; Marschner et al. 1999]. The measurements of Matusik et al. [2003] provide a dense ($90 \times 90 \times 180$) sampling of many isotropic BRDFs. The main drawback of these models is their size, since they typically represent the full 3D isotropic BRDF in tabular form. In his thesis, Matusik [2003] also describes one approach for sampling these measured BRDFs, but this representation requires as much storage as the original BRDF, making it difficult to use for scenes containing many materials.

Factored BRDF representations: In an effort to reduce the size of measured BRDF models while maintaining an accurate representation of their effects, several researchers have investigated techniques for factoring these large datasets into a more compact, manageable form [Kautz and McCool 1999; McCool et al. 2001; Suykens et al. 2003]. In all cases, the 4D BRDF is factored into products of 2-dimensional functions that can be represented as texture maps and used to shade a model in real-time. However, in most cases these factorizations allow only a single term approximation. More important, there are no techniques for importance sampling these representations.

Importance sampling: The benefit of stratified importance sampling within the context of physically-based rendering has certainly been justified by the work of Cook [1986]. Since Shirley demonstrated how to efficiently sample the traditional Phong BRDF [Shirley 1990] and Lafortune introduced a generalization of this cosine-lobe model [Lafortune et al. 1997], a reasonable approach to importance sampling an arbitrary BRDF has been to sample a best-fit approximation of one of these simpler models. Although this technique marks a clear improvement over random sampling, it has several drawbacks. First, it is not always trivial to approximate the many complex BRDFs that exist in nature with one of these models. Often, a nonlinear optimizer has difficulty fitting more than 2 lobes of a Lafortune model without careful user intervention. Second, since the sampling is only as efficient as the approximation is accurate, it is not always the case that this strategy will optimally reduce the variance for an arbitrarily complex BRDF. Our approach, on the other hand, robustly detects the energy in a BRDF during the factorization step and provides a more efficient sampling strategy.

3   A BRDF Representation for Sampling

In this section, we consider the requirements and design choices in choosing our BRDF representation for sampling. We discuss why our requirements are different from those of previous factored representations, and present a new factorization approach optimized for the needs of importance sampling. Sections 4 and 5 go on to discuss implementation details of our representation and sampling algorithms.

We begin with the observation that, in the context of a standard backward ray- or path-tracer, we will generally know the outgoing direction (θ_o, ϕ_o) and will need to sample lighting and visibility over the incident hemisphere. A straightforward approach would be to tabulate (θ_i, ϕ_i) slices of the BRDF for a dense set of directions covering (θ_o, ϕ_o), and use the appropriate one. This is essentially the approach taken by Matusik [2003] and, as far as we know, is the only previous approach for importance sampling of arbitrary measured materials. However, as already noted, this representation requires a large amount of storage space for both analytical and measured materials—as large as or larger than the tabulated BRDF itself.

Instead, we observe that for nearly all common materials there is coherence in the BRDF for different outgoing directions—for instance, the shape of the specular lobe in a glossy BRDF will often remain similar. Our goal is to exploit this coherence to develop a compact representation. First, we **reparameterize** the BRDF, e.g. by using the half-angle [Rusinkiewicz 1998]. As a number of authors have observed, reparameterization by the half-angle properly aligns BRDF features such as specular reflection, making them simpler to represent. Next, we use **factored forms**, writing the 4D BRDF as a sum of a small number of products of 2D factors, to exploit the coherence in the BRDF and develop a compact representation. Similar approaches have been developed for real-time rendering, and have shown that reparameterization and factorization can be a compact and accurate way to represent most BRDFs [Kautz and McCool 1999; McCool et al. 2001].

The specific factored decomposition we use is the following:

$$\rho(\omega_i, \omega_o)(\omega_i \cdot \mathbf{n}) \approx \sum_{j=1}^{J} F_j(\omega_o) G_j(\omega_p), \qquad (1)$$

where we have decomposed the original 4D BRDF function (multiplied by the cosine of the incident angle) into a sum of products of 2D functions. One of the functions always depends on the view direction ω_o, and the other function is dependent on some direction ω_p arising from the reparameterization. In the case of a half-angle parameterization, ω_p is taken to be

$$\omega_h = \frac{\omega_i + \omega_o}{|\omega_i + \omega_o|}. \qquad (2)$$

Unlike previous factorization approaches, this representation satisfies a number of key properties for sampling:

- **One factor dependent on outgoing direction:** When sampling according to our representation, we know the outgoing direction ω_o but not the incident direction (since we are sampling over it). Therefore, we can directly evaluate F and it is critical that it depend *only* on the outgoing direction.

- **Sum of products of two factors:** Each term above is the product of two factors F and G, where F depends only on the outgoing direction. Thus, it is easy to sample according to the second factor G only. On the other hand, approaches such as homomorphic factorization [McCool et al. 2001] or chained matrix factorization [Suykens et al. 2003] can include multiple factors in a term, making importance sampling difficult. We can also enable multiple terms (with different j) for more accurate sampling—another feature that is difficult to incorporate in homomorphic factorization.
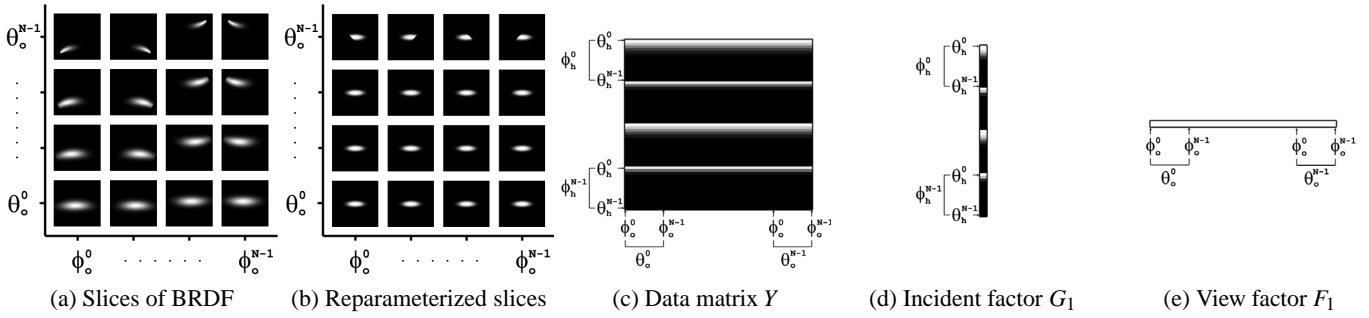
(a) Slices of BRDF	(b) Reparameterized slices	(c) Data matrix Y	(d) Incident factor G_1	(e) View factor F_1

Figure 2: *Steps in factoring a BRDF into our representation, shown for a Phong-like anisotropic BRDF [Ashikhmin and Shirley 2000]. (a) We first compute regularly-sampled 2D slices of the BRDF at a fixed set of outgoing directions. Notice that in this case the flattened highlight moves according to the position of perfect specular reflection for each view. (b) We would like to maximize the symmetry in these slices of the BRDF to make the factorization more accurate. We accomplish this by reparameterizing these values of the BRDF with respect to the half-angle vector, as the energy in this BRDF is symmetric about this direction. (c) We organize these samples into a 2D matrix by unfolding each 2D slice of the BRDF into a separate column in the matrix according to its half-angle parameterization. Notice that our choice of parameterization produces a data matrix that has a rank very close to 1. (d & e) We use non-negative matrix factorization to factor this data matrix into the outer product of two vectors. Because the rank of the original matrix was close to 1, we need only one term in the factorization. (d) In the end, we are left with a column vector dependent only on the incoming direction (G_1 reparameterized with respect to the half-angle vector in this case) and (e) a row vector F_1 dependent only on the outgoing view direction.*

- **Non-negative factors:** As opposed to using a matrix decomposition algorithm such as SVD [Kautz and McCool 1999], we use non-negative matrix factorization to ensure that all values are positive. This is necessary for interpreting the resulting factors as probability distributions, according to which we can then sample.

We now observe that most BRDFs can be compactly represented by further factoring each of the G_j into 1D functions $u_{jk}(\theta_p)$ and $v_{jk}(\phi_p)$. By doing this, we can separately treat u and v as 1D distributions, for which importance sampling is straightforward. Our final factored BRDF representation is therefore

$$\rho(\omega_i, \omega_o)(\omega_i \cdot \mathbf{n}) \approx \sum_{j=1}^{J} F_j(\omega_o) \sum_{k=1}^{K} u_{jk}(\theta_p) v_{jk}(\phi_p). \quad (3)$$

There are a total of JK terms in the final factorization, each a product of a two-dimensional function (F_j) and two one-dimensional functions (u_{jk} and v_{jk}).

It should be noted that this factorization fits the form of many isotropic and anisotropic BRDFs well. For instance, a Blinn-Phong BRDF [Blinn 1977] can be fit exactly by a single term ($J = K = 1$), with variation appearing only in the $u(\theta_h)$ function. Similarly, only 2 terms ($J = 1, K = 2$) are needed for an anisotropic Phong BRDF [Ashikhmin and Shirley 2000]. The same holds approximately for many other materials, so the above representation typically gives accurate results with a small number of terms.

Note that, unlike other uses of factored BRDF models, we typically use the representation above only for choosing samples. For actually computing the BRDF value, we use an analytic formula where available. In this case, the representation above is a compact means of sampling these commonly used analytic models, which have hitherto been difficult to sample. Similarly, for compact basis function representations, such as the Zernike polynomial expansions [Koenderink and van Doorn 1998] used in the CURET database [Dana et al. 1999] or spherical harmonic approximations [Westin et al. 1992; Sillion et al. 1991], we can use the BRDF value represented by the basis functions, using our representation only for importance sampling. In other cases, such as the dense measured BRDF representations of Matusik et al. [2003], we take advantage of the compactness of our multi-term factorization and use it as the primary representation for both reconstruction and sampling. This provides a 200-fold savings in storage in many cases, while remaining faithful to the original data.

4 Factorization

We now describe the details of our method to factor a tabular BRDF—Figure 2 provides an overview of the process. Unlike some previous methods, we use multiple non-negative terms in equations (1) and (3). This disallows common techniques such as homomorphic factorization or singular value decomposition. Instead, inspired by Chen et al. [2002], we use non-negative matrix factorization (NMF) [Lee and Seung 2000] to decompose the reparameterized BRDF. NMF is an iterative algorithm that allows multiterm factorizations, guaranteeing that all the entries in the factors are non-negative. We have found it to be robust for both single- and multi-term decompositions, and capable of producing accurate approximations for a wide range of both analytical and measured BRDFs.

Data matrix: We first organize the set of values of the original reparameterized BRDF into a matrix. We consider taking N_{θ_o} regular samples along the outgoing elevation angle and N_{ϕ_o} samples along the outgoing azimuthal angle. For each of these N_{ω_o} view directions, we record N_{ω_p} samples of the BRDF intensity (multiplied by $\cos \theta_i$), spaced equally in azimuthal and elevation angles for the chosen BRDF reparameterization. We organize the initial data samples into an $N_{\omega_p} \times N_{\omega_o}$ matrix Y.

First factorization: Using the appropriate NMF update rules, which are summarized in the appendix, we factor Y into the product of two matrices of lower dimension:

$$\begin{bmatrix} & Y & \end{bmatrix} = \begin{bmatrix} G \end{bmatrix} \begin{bmatrix} & F & \end{bmatrix} \quad (4)$$

As shown in Figure 2, G is an $N_{\omega_p} \times J$ matrix, with each column corresponding to a factor G_j in equation (1), while F is a $J \times N_{\omega_o}$ matrix, with each row corresponding to a factor F_j in equation (3). It can be helpful to interpret the G_j factors as basis images over the incoming hemisphere and F_j as encoding the appropriate mixing weights to approximate the original BRDF. In practice, we rarely need more than 3 or 4 terms to achieve an accurate approximation, and we always reduce the size of the original BRDF by at least an order of magnitude since our factored representation involves 2D functions rather than a 3D or 4D BRDF.

(a) Incident factor G_1 (reshuffled from Fig. 2d) (b) u_{jk} terms   (c) v_{jk} terms

Figure 3: *In order to optimize our representation for importance sampling, we perform another factorization step on the 2D functions dependent on the incoming direction (a) We first re-organize each column in G into a matrix such that the rows and columns vary with respect to the elevation and azimuthal angles respectively. We again apply NMF to decompose this matrix into an outer product of terms. In this example, we choose to factor this matrix into two terms. In the end, we are left with (b) two column vectors that each depend only on the elevation angle and (c) two row vectors that depend only on the azimuthal angle of the incoming direction. These are the u and v terms, respectively, in our final representation.*

If we were interested only in reducing the size of the BRDF, or in using our representation for real-time rendering, it might be reasonable to use these 2D functions directly. Notice, however, that we have approximated the intensity of the BRDF, $\rho_{int}(\omega_i, \omega_o)$ times the cosine term, which is appropriate for importance sampling. Therefore, we would need to update this representation to account for the wavelength dependence in the original BRDF. We accomplish this by using NMF to compute a single-term approximation of the BRDF at a particular wavelength (e.g. red, green or blue) divided by the intensity. For the red color channel, we would factor a data matrix composed of samples of the function: $\rho_{red}(\omega_i, \omega_o)/\rho_{int}(\omega_i, \omega_o)$ and reconstruct the red value of the BRDF by scaling our approximation of the intensity by the approximation of this function.

There also remains a challenge in sampling according to the 2D distribution G_j. It is possible to use explicit tabular approaches, by storing a Cumulative Distribution Function on θ_p for each ϕ_p, but such representations are not compact. Furthermore, effectively generating stratified samples given these 2D tabulated CDFs proves to be a difficult problem[1]. Therefore, we perform a second factorization of G_j into 1D functions dependent on θ_p and ϕ_p, which not only matches the form of most common BRDFs, but also makes the representation easy to sample and further reduces storage requirements.

Second factorization: As shown in Figure 3, we separately factor each column of the matrix G, corresponding to a 2D function that depends on the reparameterized incoming direction (θ_p, ϕ_p):

$$\left[\begin{array}{c} G_j \end{array} \right] = \left[\begin{array}{c} u_j \end{array} \right] \left[\begin{array}{c} v_j \end{array} \right], \quad (5)$$

where $u_j(\theta_p)$ is an $N_{\theta_p} \times K$ matrix, with each column corresponding to a factor u_{jk} in equation (3), and $v_j(\phi_p)$ is a $K \times N_{\phi_p}$ matrix, with each row corresponding to a factor v_{jk}.

Normalization: For the purposes of sampling, it is desirable to treat u_{jk} and v_{jk} as normalized 1D probability distribution functions. To do this, we first define

$$\overline{u}_{jk} = \int_0^\pi u_{jk}(\theta_p) \sin\theta_p \, d\theta_p, \quad \overline{v}_{jk} = \int_0^{2\pi} v_{jk}(\phi_p) \, d\phi_p. \quad (6)$$

---

[1] Effective 2D stratification in the context of environment map sampling has been proposed by Agarwal et al. [2003] using Hochbaum-Shmoys clustering, but this approach requires fixing the number of samples a priori, while in our case the number of samples for each G_j depends on the view.

Then we normalize each term T_{jk} as

$$
\begin{aligned}
T_{jk} &= F_j(\omega_o)\, u_{jk}(\theta_p)\, v_{jk}(\phi_p) \\
&= \left(\overline{u}_{jk}\, \overline{v}_{jk}\, F_j(\omega_o) \right) \left(\frac{u_{jk}(\theta_p)}{\overline{u}_{jk}} \right) \left(\frac{v_{jk}(\phi_p)}{\overline{v}_{jk}} \right) \\
&= F'_{jk}(\omega_o)\, u'_{jk}(\theta_p)\, v'_{jk}(\phi_p).
\end{aligned}
\quad (7)
$$

Finally, dropping the primes and using a single index l, we obtain the final form of our factored representation, where u_l and v_l are proper 1D probability distribution functions,

$$\rho(\omega_i, \omega_o)(\omega_i \cdot \mathbf{n}) \approx \sum_{l=1}^{L} F_l(\omega_o)\, u_l(\theta_p)\, v_l(\phi_p), \quad L = JK. \quad (8)$$

Discussion: Our representation is designed with a view to developing a sampling algorithm, and lacks two properties that are sometimes theoretically desirable. First, the terms in equation (1), whose form is essential for sampling, do not explicitly enforce reciprocity. (Of course, since we factor the product of the BRDF and the cosine term, the input function is not reciprocal to begin with.) Second, the representation is not guaranteed to be continuous—there can be a discontinuity at the pole $\theta_p = 0$ in the second factorization in equation (3). In either case, we have not observed any drawbacks in practice because of these properties and it can be seen from our results that our multiple-term fits are accurate.

5 Sampling

We now describe how to use our representation in equation (8) for importance sampling. Intuitively, each term in the approximation corresponds to a specific "lobe" of the original BRDF, and the factorization algorithm works to find the best set of lobes to approximate its overall structure. We first randomly select one of these lobes according to the energy it contributes to the overall BRDF for the current view. Next, we sample the hemisphere according to the shape of this lobe by sequentially generating an elevation and azimuthal angle according to the 1D factors u_l and v_l.

To further demonstrate this idea, consider the pair of factors that we computed to approximate the anisotropic BRDF in Figures 2 and 3. The first term creates a pair of lobes that extend along the y-axis, centered about the specular direction (Figure 4a), and the second term creates a pair of flattened lobes that extend along the x-axis (Figure 4b). We could imagine sampling the hemisphere according to just one of these terms (Figure 4c,d): using each term alone generates samples in a different region of the BRDF. If we generate samples according to both terms with equal probability, however, the aggregate effect is that we distribute samples along the anisotropic highlight (Figure 4e).

5.1 Importance Sampling

We now describe the mathematics of sampling more formally. We will be interested in evaluating the integral of the incident illumination for a fixed outgoing direction ω_o at a given pixel with location x and surface normal $\mathbf{n}$,

$$
\begin{aligned}
&\int_{\Omega_{2\pi}} L_i(x, \omega_i)\, \rho(x, \omega_i, \omega_o)(\omega_i \cdot \mathbf{n})\, d\omega_i \\
&\approx \frac{1}{n} \sum_{s=1}^{n} L_i(x, \omega_s) \left[\frac{\rho(x, \omega_s, \omega_o)(\omega_s \cdot \mathbf{n})}{\gamma_i(\omega_s \mid \omega_o)} \right].
\end{aligned}
\quad (9)
$$

The first line is simply the reflection equation—the incident lighting L_i may be evaluated iteratively or recursively for global illumination. The second line is a Monte Carlo estimator that describes the standard approach to importance sampling. It represents a weighted average of each of the samples, each divided by the probability γ_i of generating sample direction ω_s assuming that ω_o is fixed. The subscript in γ_i denotes that the probability distribution is

1				1		
0.8				0.8		
0.6		$u_1(\theta_h)$		0.6		$u_2(\theta_h)$
0.4		$v_1(\phi_h)$		0.4		$v_2(\phi_h)$
0.2				0.2		
0				0		

(a) Term 1   (b) Term 2

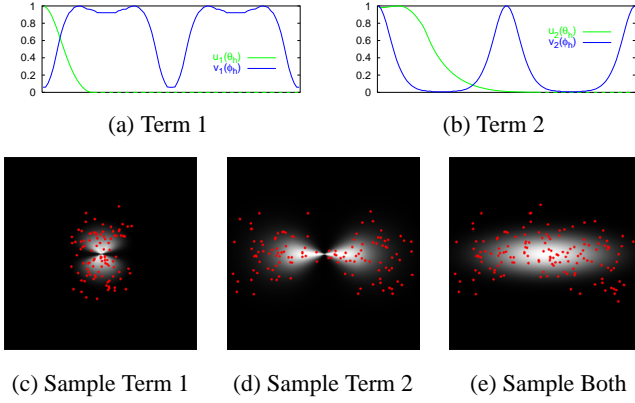(c) Sample Term 1   (d) Sample Term 2   (e) Sample Both

Figure 4: The main benefit of our factored representation is that it can be used to efficiently sample a BRDF. (a & b) In the top row, we graph the values of the 1D terms in our factorization of the BRDF considered in Figures 2 and 3. The green lines show the values of $u_1(\theta_h)$ and $u_2(\theta_h)$, whereas the blue lines represent $v_1(\phi_h)$ and $v_2(\phi_h)$. Using the strategy detailed in Section 5, we select either the first or second term to generate an incoming direction. (c) Using only the first term to generate samples, we notice that the directions accumulate around a pair of lobes along the y axis, centered within the highlight. (d) Using only the second term to generate samples, the directions accumulate around two lobes centered at $\phi_h = 0$ and $\phi_h = \pi$. (e) When we select between these two terms with equal probability, we produce a sampling pattern that matches the energy in the original BRDF.

over incident directions. γ_i should be non-negative and normalized, i.e. $\int_\Omega \gamma_i(\omega_i \mid \omega_o)\, d\omega_i = 1$. Our representation is used to generate samples ω_s, according to the probabilities γ_i. For analytic models, the actual BRDF can be used to evaluate ρ. The more accurate our representation is, the lower the variance, but equation (9) is always accurate and unbiased.

If our factored representation exactly represents the BRDF multiplied by the cosine term, the numerator in the bracketed term in equation (9) will be exactly proportional to the denominator, and that term will simply be a constant. The estimator will then represent the ideal importance sampling method based on the BRDF, and will have low variance. In fact, in the limiting case of a constant environment (L_i is constant), there will be zero variance. In practice our representation is not exact, but it is a good approximation and importance sampling with it significantly reduces variance.

5.2 Sampling Algorithm

We now describe how to choose directions ω_s and evaluate γ_i in equation (9). Our method chooses the lobe l, azimuthal angle ϕ_p, and elevation angle θ_p in turn, with each step involving computing one random number and inverting a 1D Cumulative Distribution Function.

Choosing a term l: The probability of choosing a term l, for a given outgoing direction ω_o is given by

$$\gamma(l|\omega_o) = \frac{F_l(\omega_o)}{\sum_{j=1}^{L} F_j(\omega_o)}. \tag{10}$$

From these probabilities, we calculate a 1D CDF over l. To select a term, we generate a uniform random variable in $[0, 1]$ and perform a binary search on the CDF to transform the random variable into a value of l. Notice that the probabilities depend on the view direction, so we must recompute this CDF each time the outgoing direction changes. However, L is typically very small, and the same CDF can be used for all samples through a given pixel (since ω_o is fixed), so the computation is inexpensive.

Choosing azimuthal and elevation angles ϕ_p and θ_p: Having chosen the term l to sample, we must now choose ϕ_p based on the probability distribution $v_l(\phi_p)$. As before, we generate a uniformly distributed random variable in $[0, 1]$, and numerically invert the CDF at that value to determine ϕ_p. Choosing θ_p follows the same methodology, but because of the $\sin\theta_p$ area measure, we find it simpler to define $z = \cos\theta_p$, and use $u_l(z)$ as the probability distribution. Inverting the CDF then yields z, from which we find $\theta_p = \cos^{-1} z$. Note that we can precompute these CDFs because the probabilities do not depend on ω_o—a significant benefit of our factorization.

Computing probability: Given θ_p and ϕ_p, it is straightforward to generate a direction ω_s. Due to reparameterization, it is occasionally possible for the sample directions ω_s to take values below the horizon, but we can simply set the estimator to 0 for those directions without introducing inaccuracy or bias. Otherwise, we calculate the probability for equation (9) as the sum of the marginal probabilities for each term:

$$\gamma_p(z, \phi_p | \omega_o) = \sum_{l=1}^{L} \frac{F_l(\omega_o)\, u_l(z)\, v_l(\phi_p)}{\sum_{j=1}^{L} F_j(\omega_o)}. \tag{11}$$

One issue we must address is reparameterization, since equation (9) is in terms of the incident direction ω_i while our factors are reparameterized using the half-angle or, in general, some alternative parameterization ω_p. Since it is easy to convert between them, there is no difficulty in evaluating equation (9). However, our probability distributions γ_p are in terms of the new parameterization, and must be modified to conform to equation (9). In particular,

$$\gamma_i(\omega_i \mid \omega_o) = \gamma_p(\omega_p \mid \omega_o) \left| \frac{\partial \omega_p}{\partial \omega_i} \right|, \tag{12}$$

where the last term is equivalent to the Jacobian for changing variables in multidimensional integration, and converts differential areas in ω_p to those in ω_i. For the half-angle, this function has been computed in many derivations, such as for calculating the Torrance-Sparrow BRDF [1967], and is given by

$$\left| \frac{\partial \omega_h}{\partial \omega_i} \right| = \frac{1}{4(\omega_i \cdot \omega_h)}. \tag{13}$$

Stratification: The preceding algorithm generates single samples independently, but it is straightforward to extend it to generate stratified samples: we simply stratify each of the individual stages. Because these stages depend only on 1D probability distribution functions, this is accomplished by stratifying the domain of the uniform random variables used in those stages. We have found this to be an effective method of further reducing variance in the generated images.

6 Results

We now present the results of factoring both analytical and measured BRDFs, describing the accuracy and compactness of our representation. In 6.2, we analyze the efficiency of sampling according to this representation.

6.1 Factorization

We factored four analytic BRDF models of varied behavior: the Cook-Torrance [1982] rough-surface model, an anisotropic Ward model [1992], Poulin-Fournier [1990] anisotropic reflection from cylinders, and the Hapke-Lommel BRDF [1963] with strong backscattering effects. We also tested three measured BRDFs acquired by Matusik et al. [2003]: nickel, plastic and metallic blue. Table 1 lists the resolution and parameterization of each factorization along with the normalized mean absolute error (MAE) in the approximation. These errors were computed over a dense set of samples of

Original BRDF	Resolution $(\theta_o \times \phi_o \times \theta_p \times \phi_p)$	Terms $(J \times K = L)$	Param.	Compression Ratio	Normalized MAE Factored	Normalized MAE Lafortune
Cook-Torrance	$16 \times 16 \times 32 \times 16$	$4 \times 1 = 4$	ω_h	n/a	0.192	0.632
Ward	$16 \times 16 \times 100 \times 100$	$2 \times 4 = 8$	ω_h	n/a	0.094	1.092
Poulin-Fournier	$16 \times 16 \times 32 \times 16$	$3 \times 1 = 3$	ω_i	n/a	0.142	0.348
Hapke-Lommel	$16 \times 16 \times 32 \times 16$	$3 \times 1 = 3$	ω_i	n/a	0.186	0.464
Measured Nickel	$16 \times 16 \times 128 \times 16$	$2 \times 1 = 2$	ω_h	230:1	0.201	0.643
Measured Plastic	$16 \times 16 \times 128 \times 16$	$3 \times 1 = 3$	ω_h	200:1	0.266	0.874
Measured Metallic-Blue	$16 \times 16 \times 128 \times 16$	$4 \times 1 = 4$	ω_h	180:1	0.118	0.464

Table 1: *Accuracy of the factored BRDF representation. We factored 4 analytical BRDFs: Cook-Torrance ($d = 0.1, R_d = [0.12, 0.22, 0.48], s = 0.9, F_0 = R_d, m = 0.2$), Ward ($\rho_d = 0.1, \rho_s = 1.2, \alpha_x = 0.2, \alpha_y = 0.02$), Poulin-Fournier ($d = 2.0, h = 0.0, n = 20.0, R_s = 0.8, R_d = 0.2$), and Hapke-Lommel ($g = 0.6, f = 0.1, r = 1.0$), along with 3 measured BRDFs from Matusik et. al. [2003]: nickel, plastic, and metallic-blue. For each BRDF we list the resolution of the original data matrix Y, the number of "outer" and "inner" terms (J and K, respectively) of the factorization, and the parameterization of the incoming hemisphere. For the measured BRDFs, we also list the compression ratio. We report the mean absolute error of the final factorization, normalized by the mean BRDF value. This is compared to the error resulting from fitting a multi-lobe Lafortune model to the original BRDF using a standard non-linear optimizer.*

the 4D domain, independent of the resolution of the factorization. We compare this with a best-fit 2-lobe Lafortune model, except for the Cook-Torrance BRDF, to which we fit a 3-lobe model. For most models, we reparameterized by the half-angle $\omega_p = \omega_h$, while for the more diffuse models (Poulin-Fournier and Hapke-Lommel), we used the standard parameterization by incident angle $\omega_p = \omega_i$. Compression ratios are reported for measured BRDFs, and represent the reduction in size with respect to the original data matrix.

We see that in all cases factorization produces an accurate result, in many cases significantly more accurate than fitting an analytic model such as Lafortune. This accuracy in the representation explains the high quality of our sampling algorithm. We further note that fitting a 3-lobe Lafortune model can be unstable, often taking minutes to hours to converge in a nonlinear minimizer, and can require manual tuning to find a good fit. By contrast, our method is automatic, robust, and fast (taking only a few minutes to factor the BRDFs considered in these experiments).

We observe, as previous authors have, that MAE or RMS errors are imperfect measures of the accuracy and visual quality of a BRDF approximation: in practice, the numerical error is dominated by regions such as the specular highlight and grazing angles. To this end, Figure 5 shows the appearance of some factored models, as compared to the originals, under point illumination. We see that throughout most of the BRDF the representation accuracy is, in fact, better than the numbers in Table 1 would suggest, and the error of our approximation decreases rapidly as more terms are added (Figure 6). For the case of measured nickel, note that our representation regularizes some of the measurement noise around the highlight, relative to the original data. We conclude that, for measured data, our representation appears to produce results comparable with measurement error (Matusik et al. observe errors, such as deviation from reciprocity, of $10 - 15\%$ at normal angles, ranging to $60 - 70\%$ at grazing angles [personal communication]).

Selecting the appropriate resolution for the factorization, and the parameterization of the incoming hemisphere is a manual process. In most cases, the analytical formula (for parametric BRDFs) or general appearance (for measured BRDFs) provides enough information for an accurate estimate of how many samples are sufficient and what parameterization is optimal. Theoretically, the number of terms should be proportional to the ranks of the matrices Y and G_j (or, at least, the number of significant eigenvalues of these matrices). In practice, however, we simply increase the number of terms (J and K) until the error in the approximation plateaus. Figure 6 shows this convergence process for factorizations of the anisotropic Ward BRDF listed in Table 1.

Since our goal is to develop a representation suitable for efficient sampling, rather than a factorization method more accurate than previous approaches, we did not directly compare with previ-
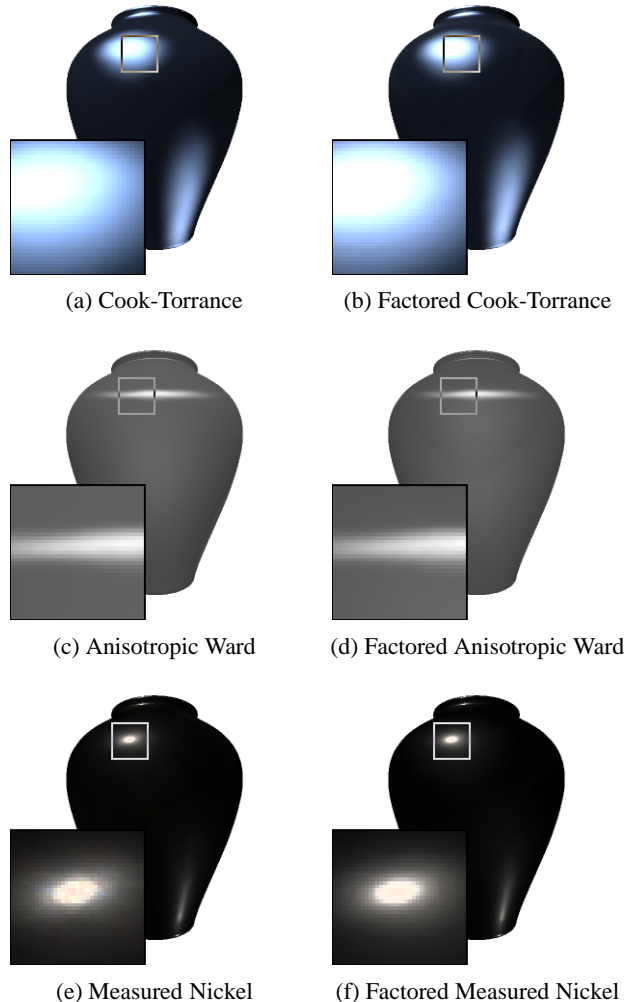

(a) Cook-Torrance (b) Factored Cook-Torrance

(c) Anisotropic Ward (d) Factored Anisotropic Ward

(e) Measured Nickel (f) Factored Measured Nickel

Figure 5: *Accuracy of the BRDF factorization. The left column shows a vase rendered with (a) a Cook-Torrance BRDF, (c) a Ward anisotropic BRDF and (e) a measured nickel BRDF under direct illumination. (b, d & f) The right column shows the same vase rendered with a factored approximation of the original BRDF. (b) Notice the slight banding effects that appear in the factored highlight of the Cook-Torrance BRDF, which result from the finite sampling resolution along θ_h. (f) The factorization actually regularizes some of the measurement noise that appears in the highlight of the measured nickel BRDF.*
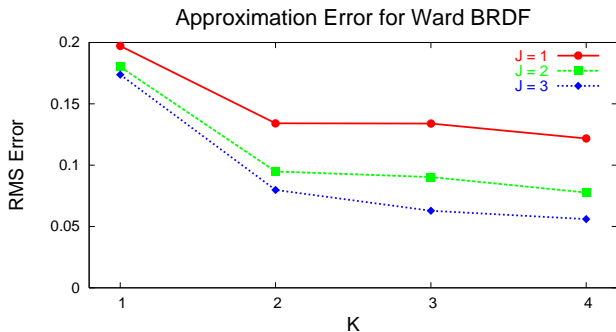
Figure 6: *Comparison of the RMS approximation error of an anisotropic Ward BRDF, as a function of the number of terms in the factorization. Each line shows a different number of "outer" terms (J) while the number of "inner" terms (K) increases along the x axis. Note the drop in error at K = 2: this shows that at least two inner terms are necessary to capture the anisotropic shape of this BRDF.*

ous factorization approaches that cannot be easily sampled. However, we did factor a Poulin-Fournier model with qualitatively comparable parameters to the one listed in [McCool et al. 2001], and produced a factorization with RMS error comparable with that approach (although not given in Table 1, the RMS error for that factorization is 0.094). While this is not the focus of our paper, these results indicate that the benefits of a multi-term nonnegative factorization may be applicable in other areas such as real-time rendering.

6.2 Sampling

We next consider the efficiency of importance sampling using our factored representation. For a controlled quantitative comparison, we conducted tests involving images of a sphere (so visibility is not considered), lit by a constant environment map (so complex illumination is not considered). The comparison methods are uniform sampling of a cosine-weighted hemisphere, analytic sampling of either a best-fit multi-lobe Lafortune model [Lafortune et al. 1997] or a generalized Blinn-Phong model developed by Ashikhmin and Shirley [2000], and an approach based on explicit tabulation of the BRDF [Matusik 2003]. All methods were stratified.

We compared variance (averaged over 50 trials) as a function of the number of samples used (ground truth was taken as the limit with a very large number of samples) for the BRDFs considered in Table 1. We verified for all sampling techniques that they were unbiased, and that the image variance decayed approximately as the inverse of the number of samples (Figure 7). Table 2 reports the ratio of the variance of the comparison methods to our approach with 100 samples—the relative performance with a different sample count would be essentially the same. This is an appropriate metric, since it directly corresponds to how much longer the alternative approaches would need to run (i.e., how many more samples they would require) to produce the same quality results as our method. The image RMS error corresponds roughly to the standard deviation, which is the square root of the variance.

We see that compared to uniform random sampling, BRDF importance sampling always does at least 5 to 10 times better, and significantly better for shiny materials such as measured nickel. Relative to analytic models, the degree of improvement depends on how closely the analytic model is able to match the BRDF. Lafortune's model, for instance, is a good fit of the Poulin-Fournier and Hapke-Lommel BRDFs (as seen in Table 1). Note that these materials are more diffuse and random sampling also does fairly well on them. However, we always do at least twice as well as sampling based on a Lafortune fit, and for measured materials, and even for the widely known Cook-Torrance model, we do an order of magnitude better. On the other hand, the Ashikhmin-Shirley model represents


Figure 7: *Image variance as a function of the number of samples. These plots show the relationship between the average image variance and the number of samples/pixel for all 5 sampling strategies considered in this paper.* **Top:** *the variance in the image of a sphere with the Cook-Torrance BRDF from Table 1 under constant illumination.* **Bottom:** *variance in the image of a sphere with a measured metallic-blue BRDF under constant illumination. As expected, the variance converges to 0 as the sample counts increase, confirming that each strategy produces an unbiased estimate. At 100 paths/pixel we see the values for which the factor of improvement is listed in Table 2.*

lobes depending on the half-angle well, and therefore does better than Lafortune at sampling metals and plastics such as the Cook-Torrance, nickel and plastic BRDFs. However, our method is still at least a factor of 2 better, and for many of the materials, we see an improvement by a factor of 5-10.

We also measured the effectiveness of importance sampling the BRDF using the factored representation under complex illumination (Table 3). The experimental setup is identical to that for Table 2, except that the sphere was placed in the beach environment [http://www.debevec.org/Probes/]. Because the shape of the integrand of the rendering equation is affected by varying illumination, sampling the BRDF alone will not be as efficient as for constant illumination. As expected, for more diffuse BRDFs (Hapke-Lommel and Poulin-Fournier) we notice that the illumination becomes the dominant factor in the integrand and uniform random sampling is a reasonable strategy. For the more specular BRDFs, however, we still see the benefits of importance sampling the BRDF, and our method decreases the variance by a factor of 2-20 over best-fits of either parametric formula. One example of this is the measured metallic-blue BRDF. The specular peak of this BRDF deviates substantially from the ideal specular direction, and is also not well approximated by a function of θ_h, particularly as the view approaches the horizon. As a result, the best-fit Lafortune and Ashikhmin-Shirley parametric models fail to match the BRDF well in these regions. Although our factored representation uses a half-angle parameterization of the incoming hemisphere as well, it can handle small deviations from this direction through the inherent flexibility a numerical approximation provides. As a result, our technique samples this BRDF more efficiently than either parametric fit (Fig-

Efficiency of BRDF Sampling in Constant Environment

Original BRDF	Improvement relative to:			
	Unif.	Laf.	A&S	Mat.
Cook-Torrance	16.38	13.27	3.53	0.75
Poulin-Fournier	5.86	1.85	6.11	n/a
Hapke-Lommel	3.32	2.14	11.61	1.99
Nickel	306.17	11.52	2.17	1.66
Plastic	157.12	14.40	1.34	18.53
Met. Blue	8.88	6.73	6.75	0.44

Table 2: *Efficiency of importance sampling the BRDF. This table lists the factor of improvement in variance resulting from sampling the BRDF according to our factored representation, compared to four alternative approaches: uniformly sampling a cosine-weighted hemisphere, sampling a best-fit multi-lobe Lafortune model, sampling a best-fit generalized Blinn-Phong model described by Ashikhmin and Shirley, and sampling from a dense set of tabulated CDFs, as described by Matusik. Because variance is linearly proportional to running time, these values can be interpreted as the factor of time, or number of paths, that would be required for the other sampling approaches to reach the same noise level as our representation.*

Efficiency of BRDF Sampling in Beach Environment

Original BRDF	Improvement relative to:			
	Unif.	Laf.	A&S	Mat.
Cook-Torrance	25.79	10.28	2.23	0.75
Poulin-Fournier	1.40	1.09	1.53	n/a
Hapke-Lommel	0.89	1.61	1.87	1.00
Nickel	572.76	3.45	2.17	4.80
Plastic	381.94	21.60	1.67	55.64
Met. Blue	9.17	5.96	5.71	0.55

Table 3: *Importance sampling the BRDF under complex illumination. This table presents the factor of improvement of our sampling strategy compared to alternative approaches when rendering a particular BRDF in the beach environment. Because the illumination contributes to the shape of the integrand in the rendering equation, sampling according to the BRDF alone will be less efficient than when the illumination is constant. Although our factored representation still outperforms the alternative sampling strategies by a factor of 2-20, these results suggest the potential desirability of combining environment and BRDF sampling.*

ure 8). Together, these results indicate the generality and efficacy of our approach for importance sampling compared to fitting a specific analytic model.

The only method competitive with ours is that of Matusik [2003]. For a set of fixed view directions, this method computes a 2D CDF over incident directions according to the spherical parameterization of the hemisphere. For accurate results, this approach requires dense sampling along all variables, and does not provide the compactness of our factored representation. In fact, Matusik reports using resolutions of $90 \times 90 \times 180$ for isotropic materials, and acknowledges the infeasibility of this approach for anisotropic materials. Even with these resolutions, however, there are still situations when the closest CDF (i.e. the closest view for which the CDF is tabulated) differs significantly from the actual shape of the BRDF. This is apparent with measured nickel and measured plastic, for which the BRDF has a sharp specular peak. For views near normal incidence, sampling according to the spherical coordinates of the incident direction is sufficient to accurately capture the shape of the BRDF. Near grazing angles, however, the 2D CDF for the nearest view often varies significantly from the actual shape of the BRDF, degrading the sampling efficiency in these regions (Figure 9). Our factored representation, on the other hand, avoids this situation through a better parameterization of the hemisphere and a more continuous approximation of the BRDF over all views. Moreover, our representation supports anisotropic reflection and is more compact. For the BRDFs presented in this paper, the complete
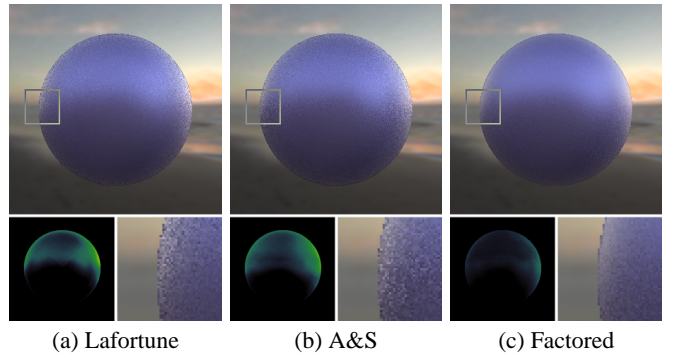

(a) Lafortune   (b) A&S   (c) Factored

Figure 8: *Importance sampling a BRDF according to best-fit parametric models and our factored representation (cf. last row of Table 3). These images show a metallic-blue sphere in the beach environment, rendered with 100 samples generated according to (a) a best-fit 2-lobe Lafortune model, (b) a best-fit Ashikhmin-Shirley model, and (c) our factored representation. We show both a variance plot on a logarithmic scale and a closeup at a region where the view approaches the horizon. In this part of its domain, the BRDF has a shape that is difficult to fit with either of the parametric models, and our factored representation allows more efficient sampling.*


Matusik Sampling   Factored Sampling

Figure 9: *Sampling measured nickel with a dense set of 2D CDFs, as described by Matusik et. al. and using our factored representation (cf. fourth row of Table 3). For such shiny BRDFs, computing a fixed set of 2D CDFs can still cause problems for regions of the domain for which the nearest precomputed CDF of a particular view poorly matches the actual BRDF. Our factored representation, on the other hand, gains better continuity through an appropriate parameterization and approximation, resulting in more efficient importance sampling throughout the domain.*

factored representation requires roughly 200KB as compared to the 60MB required to store the samples of the 3D BRDF along with the pre-computed 2D CDFs required for the approach of Matusik.

In generating a sample using our approach, the dominant cost is that of inverting three 1D CDFs using a binary search. This makes our approach reasonably fast, comparable with analytically drawing a sample according to the Lafortune and Phong sampling algorithms. It is somewhat slower than the simpler random sampling, and almost identical to the approach of Matusik, which also inverts a pair of 1D CDFs. In practice, all of these times are small compared to the cost of propagating a sample or tracing a ray for global illumination, and hence the number of samples (and the results in Table 2) corresponds closely to actual running time.

6.3 Global Illumination

We also rendered a complex scene with global illumination using a path tracer (Figure 10). In this case, the incident illumination and visibility are unknown and, consequently, importance sampling the BRDF is the only reasonable strategy (i.e., environment sampling is not possible). We used our factored representation to sample all five BRDFs in the scene and to represent the three measured BRDFs. We compare our results with those of a system using best-fit Lafortune models to sample the different BRDFs. We present rendered images at equal time (300 paths/pixel for both sampling strategies) and equal quality (1200 paths/pixel for Lafortune sampling) along with false-color visualizations of the variance in the scene on a logarithmic scale and several magnified views showing different BRDFs in the scene. Clearly, different regions of the scene converge at different rates, but our method is roughly 4-5 times more efficient overall and an order of magnitude more efficient on difficult BRDFs such as the plastic handle. This example highlights the usefulness of a general approach to both representing and importance sampling BRDFs.

7 Conclusions and Future Work

This work addresses a long-standing graphics problem of efficiently importance sampling complex analytic and measured BRDFs. We introduce a new factored representation of the BRDF that reduces sampling to inverting three 1D cumulative distribution functions. This provides a compact practical representation and a simple algorithm for sampling, which in many cases reduces variance and sampling times relative to previous methods. We use our representation and importance sampling method to render scenes with multiple isotropic and anisotropic materials with global illumination and shadows.

In future work, we would like to extend our technique to allow for mixed parameterizations of the factored BRDF, such that each term may have a different parameterization. This would allow us to better approximate BRDFs that exhibit several different types of scattering (e.g. side, backward and forward) at the same time. A second area of future research is investigating how to combine strategies for sampling the BRDF and the incident illumination. In general, illumination sampling and BRDF sampling are complementary techniques, and it would be interesting to investigate multiple importance sampling methods [Veach and Guibas 1995] for combining our algorithm with environment map sampling.

Our factorization method might also have applications in sampling bi-directional texture functions (BTFs) and light fields—two examples of high-dimensional functions that, like BRDFs, typically have significant redundancy. More generally, we see our work as a first step towards efficient techniques to sample high-dimensional measured functions. With the increasing importance of measured and image-based data in computer graphics, this problem promises to have growing significance.

8 Acknowledgements

References

AGARWAL, S., RAMAMOORTHI, R., BELONGIE, S., AND JENSEN, H. W. 2003. Structured importance sampling of environment maps. In *SIGGRAPH 03*, 605–612.

ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic Phong BRDF model. *Journal of Graphics Tools: JGT 5*, 2, 25–32.

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *SIGGRAPH 77*, 192–198.

CHEN, W.-C., BOUGUET, J.-Y., CHU, M. H., AND GRZESZCZUK, R. 2002. Light field mapping: efficient representation and hardware rendering of surface light fields. In *SIGGRAPH 02*, 447–456.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1, 7–24.

COOK, R. L. 1986. Stochastic sampling in computer graphics. *ACM Transactions on Graphics 5*, 1, 51–72.

DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1, 1–34.

GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P., ARVO, J., LAFORTUNE, E., FERWERDA, J. A., WALTER, B., TRUMBORE, B., PATTANAIK, S., AND FOO, S.-C. 1997. A framework for realistic image synthesis. In *SIGGRAPH 97*, 477–494.

HAPKE, B. 1963. A theoretical photometric function for the lunar surface. *Journal of Geophysical Research 68*, 15.

HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. In *SIGGRAPH 91*, 175–186.

KAJIYA, J. T. 1985. Anisotropic reflection models. In *SIGGRAPH 85*, 15–21.

KAJIYA, J. T. 1986. The rendering equation. In *SIGGRAPH 86*, 143–150.

KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary BRDFs using separable approximations. In *Proceedings of the 10th Eurographics Workshop on Rendering*, 281–292.

KOENDERINK, J., AND VAN DOORN, A. 1998. Phenomenological description of bidirectional surface reflection. *JOSA A 15*, 11, 2903–2912.

KOLLIG, T., AND KELLER, A. 2003. Efficient illumination by high dynamic range images. In *Eurographics Symposium on Rendering 03*, 45–51.

LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *SIGGRAPH 97*, 117–126.

LEE, D. D., AND SEUNG, H. S. 2000. Algorithms for non-negative matrix factorization. In *NIPS*, 556–562.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 1999. Image-based BRDF measurement including human skin. In *Proceedings of 10th Eurographics Workshop on Rendering*, 139–152.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. In *SIGGRAPH 03*, 759–769.

MATUSIK, W. 2003. *A Data-Drive Reflectance Model*. PhD thesis, Massachusettes Institute of Technology.

MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH 01*, 185–194.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometric Considerations and Nomenclature for Reflectance*. National Bureau of Standards (US).

OREN, M., AND NAYAR, S. K. 1994. Generalization of Lambert's reflectance model. In *SIGGRAPH 94*, 239–246.

PHONG, B. T. 1975. Illumination for computer generated pictures. *Commun. ACM 18*, 6, 311–317.

POULIN, P., AND FOURNIER, A. 1990. A model for anisotropic reflection. In *SIGGRAPH 90*, 273–282.

RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop '98*, 11–22.
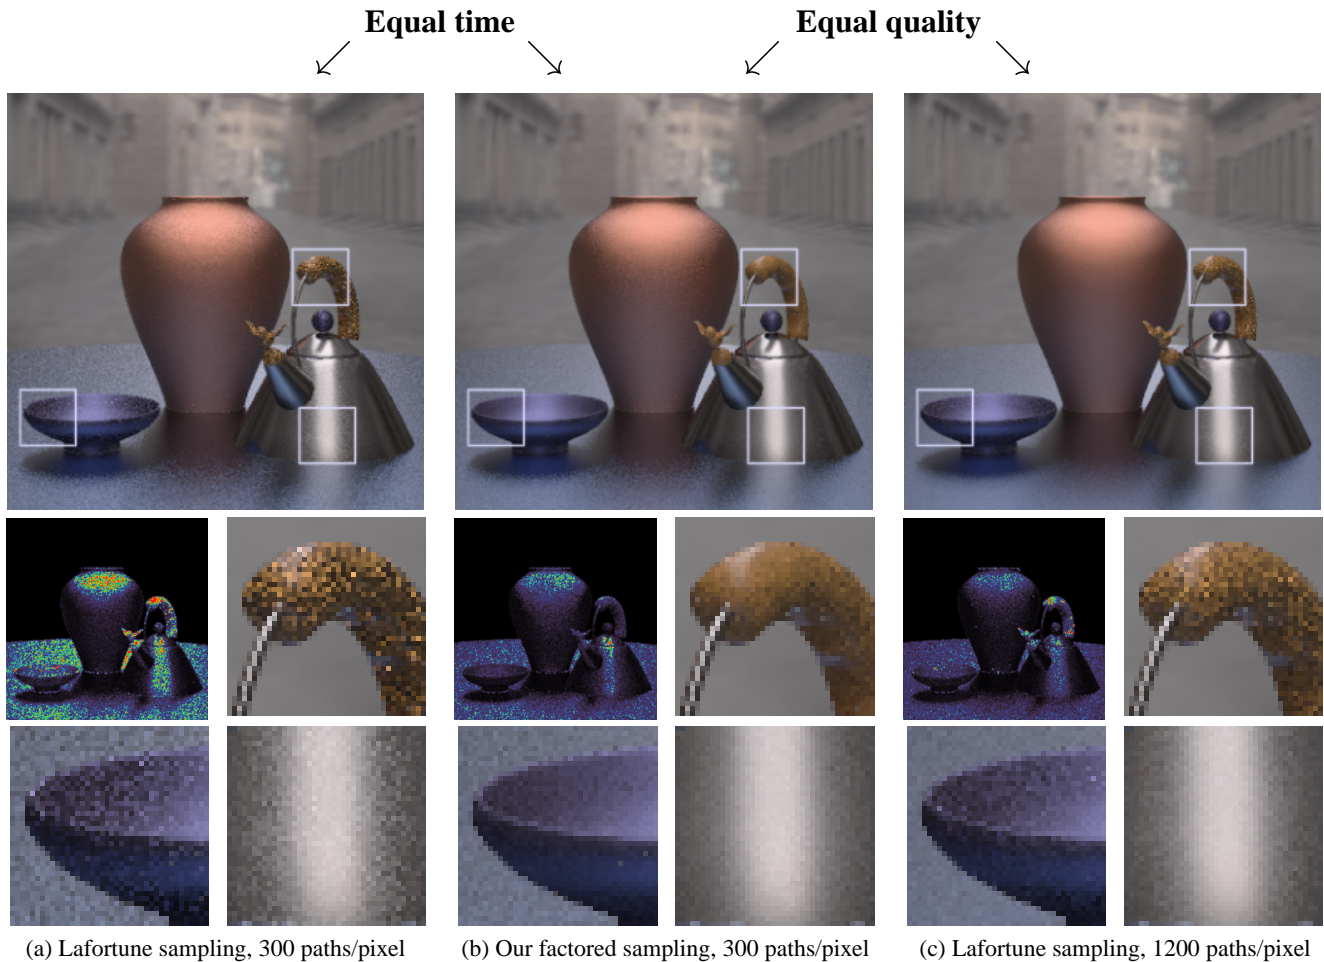
Equal time   **Equal quality**

(a) Lafortune sampling, 300 paths/pixel   (b) Our factored sampling, 300 paths/pixel   (c) Lafortune sampling, 1200 paths/pixel

Figure 10: *This scene was rendered using a path tracer for global illumination. (a) We generated 300 stratified importance samples of the local hemisphere using a best-fit Lafortune model for each of the 5 BRDFs in the scene. (b) We sampled the incoming hemisphere using our factored representation of each BRDF. (c) Lafortune sampling with 1200 samples. The bottom row shows a false-color variance plot and closeups of some regions. On the whole, we see that our method is approximately four times more efficient than Lafortune sampling, and substantially better for difficult BRDFs such as the plastic teapot handle. We refer the reader to [http://www.cs.princeton.edu/gfx/proj/brdf/fig10] for the data and algorithms developed in this paper that were used to generate this comparison.*

SHIRLEY, P. 1990. *Physically Based Lighting Calculations for Computer Graphics.* PhD thesis, University of Illinois at Urbana Champaign.

SILLION, F. X., ARVO, J. R., WESTIN, S. H., AND GREENBERG, D. P. 1991. A global illumination solution for general reflectance distributions. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, ACM Press, 187–196.

SUYKENS, F., VOM BERGE, K., LAGAE, A., AND DUTRE, P. 2003. Inter-active rendering with bidirectional texture functions. *EUROGRAPHICS 2003, Computer Graphics Forum 22*, 3.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America 57*.

VEACH, E., AND GUIBAS, L. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *SIGGRAPH 95*, 419–428.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, 265–272.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, ACM Press, 255–264.

Appendix: Non-negative Matrix Factorization

We use non-negative matrix factorization (NMF) [Lee and Seung 2000] to decompose a matrix of BRDF values, Y, into two lower-dimensional factors $Y \approx GF$. NMF is an iterative algorithm and initially both G and F are seeded with non-negative random values. At each iteration, the entries in the current factors are updated according to the deviation between their current approximation of the target matrix, (GF), and the actual target matrix Y using the following update rules:

$$F_{ij} \leftarrow F_{ij} \sum_k G_{ki} \frac{Y_{kj}}{(GF)_{kj}}, \tag{14}$$

$$G_{ij} \leftarrow G_{ij} \sum_k \frac{Y_{ik}}{(GF)_{ik}} F_{jk}, \tag{15}$$

$$G_{ij} \leftarrow \frac{G_{ij}}{\sum_k G_{kj}}. \tag{16}$$

This algorithm is guaranteed to converge to a *local* minimum of the following error metric:

$$D(Y\|GF) = \sum_{ij}\left(Y_{ij}\log\frac{Y_{ij}}{(GF)_{ij}} - Y_{ij} + (GF)_{ij}\right). \tag{17}$$

Because for a BRDF relative difference is more perceptible than absolute difference, we found that minimizing this "divergence" error was desirable.