

# A New Change of Variables for Efficient BRDF Representation

Szymon M. Rusinkiewicz  
Stanford University  
Gates Building, Wing 3B  
Stanford, CA 94305  
smr@cs.stanford.edu

## Abstract

We describe an idea for making decomposition of Bidirectional Reflectance Distribution Functions into basis functions more efficient, by performing a change-of-variables transformation on the BRDFs. In particular, we propose a reparameterization of the BRDF as a function of the halfangle (i.e. the angle halfway between the directions of incidence and reflection) and a difference angle instead of the usual parameterization in terms of angles of incidence and reflection. Because features in common BRDFs, including specular and retroreflective peaks, are aligned with the transformed coordinate axes, the change of basis reduces storage requirements for a large class of BRDFs. We present results derived from analytic BRDFs and measured data.

## 1 Introduction - BRDF Representation for Computer Graphics

Historically, the reflection models used by computer graphics renderers have been limited. Despite their physical inaccuracy, simple equations such as the Phong lighting model remain popular. True photorealism, however, will require more sophisticated and accurate models of surface properties.

The major difficulty in moving to more sophisticated reflection models has been the difficulty in representing these BRDFs efficiently. The domain of computer graphics requires BRDF representations that are accurate, have high angular resolution, and cover the entire range of possible angles of incidence and reflection. Traditionally, computer graphics systems either relied on analytic models (which were not always available for the exact surface that had to be represented) or had to store enormous quantities of data to represent even relatively simple, smooth BRDFs.

We will examine some of the approaches taken to storing BRDFs in the past, including both analytic models and decompositions into basis functions. We then present an approach for reducing the number of basis functions required to represent a BRDF by reparameterizing the BRDF in terms of the halfangle and a difference angle, rather than the usual angles of incidence and reflection. We show the savings in storage achieved by this transformation for several classes of commonly-encountered BRDFs, including BRDFs with specular, retroreflective, and anisotropic peaks. Note that in this paper we discuss only the directional dependence of BRDFs, not variation with wavelength.

## 2 Previous Approaches to BRDF Representation

The main efforts in BRDF representation have focused on either analytic formulas that can represent some narrow class of BRDFs, or generic techniques suitable for storing arbitrary four-dimensional functions. We review some of the principal work in both areas.

## 2.1 Analytic Models - Physically-based and Phenomenological

Most renderers today use BRDFs computed by an analytic formula. Many of these formulas are the result of modeling the properties of a real surface and mathematically computing the amount of light that would be reflected by a surface with those properties. For example, physically-based BRDFs have been derived for primarily specular surfaces (e.g. the Cook-Torrance-Sparrow model [Cook 81]), for rough diffuse surfaces (the Oren-Nayar model [Oren 94]), and for dusty surfaces (the Hapke/Lommel-Seeliger model, developed to model lunar reflectance [Hapke 63]). They range in complexity from simple formulas that ignore many real-world effects to complex models that attempt to account for most actually-observed surface phenomena (e.g. the He-Torrance-Sillion-Greenberg model [He 91]). Because they are derived from physical principles, these models, to a large degree, satisfy the criteria of physical plausibility, such as energy conservation and Helmholtz reciprocity.


An inherent property of physically-based models is that because they start with specific assumptions about microgeometry, they can only predict the reflectance of surfaces that closely match those assumptions. As a result, even the most elaborate theoretical models cannot predict reflectance from surfaces with complex microstructure or mesostructure (e.g. cloth, metallic paint, fur). Phenomenological models have been used successfully to widen the range of representable BRDFs. For example, the Minnaert BRDF [Minnaert 41] was an early empirical formula developed to characterize the reflectance of the moon before more physically-correct models were derived. Lafortune's generalized cosine lobes [Lafortune 97] can represent a wide range of phenomena, including off-specular reflection and retroreflection. As with all phenomenological models, however, there exist behaviors that Lafortune's functions cannot represent (including most commonly-seen kinds of anisotropy). Shade trees [Cook 84] and shading languages such as RenderMan [Hanrahan 90] attempt to generalize phenomenological models even further, by allowing simpler models to be combined in flexible ways.

As we have seen, both physically-based and phenomenological models can only represent certain limited classes of surfaces. Given an arbitrary BRDF, whether it was measured directly or obtained through simulation (as in [Westin 92]), there is no guarantee that any analytic model can represent it. This often does not meet the requirements of computer graphics, since in general one would like to produce realistic renderings of arbitrary surfaces. Thus, despite the simplicity and utility of analytic models, there have been several attempts to look at BRDFs in more general frameworks, in which it is possible to represent exactly any given BRDF.

## 2.2 Decomposition into Basis Functions

Describing a complex function as a linear combination of some set of basis functions is a widely used technique for representing continuous functions. In the domain of BRDF representation, the most popular classes of basis functions are tensor products of the spherical harmonics, Zernike polynomials, and spherical wavelets.

Spherical harmonics, the spherical analogue of sines and cosines, are a popular choice for representing BRDFs. Because spherical harmonics are compact in frequency space, smooth BRDFs should have fewer nonzero (or at least non-negligible) coefficients than complex ones when expressed in this basis. Westin et al. present an implementation of using spherical harmonics to represent BRDFs, taking advantage of symmetry and reciprocity in the BRDF to reduce storage requirements [Westin 92].



**Figure 1:** Ringing caused by truncation of high-frequency terms. The graphs are gonio-metric plots of the BRDF as a function of reflectance angle, for a fixed angle of incidence. The directions of incident light and ideal specular reflection are also shown.

An alternative to using a basis of spherical harmonics, which treat functions on a hemisphere as a special case of functions on a sphere, is to use a basis of functions on a disk, then map those functions onto a hemisphere. Koenderink et al. explore this possibility, looking at representing BRDFs in terms of the Zernike polynomials, which form an orthonormal basis of functions on the unit disk [Koenderink 96]. The authors use an equal-area mapping from the disk onto the hemisphere, and enforce reciprocity by taking particular linear combinations of the functions. Thus, the paper develops a representation very similar to that used by Westin et al., but optimized for the hemisphere rather than the sphere.

Both the spherical harmonics and Zernike polynomials require large numbers of basis functions in order to represent quickly-varying BRDFs accurately. The consequence of using too few terms is “ringing” in the BRDF, caused by sharp edges in frequency space (see Figure 1). Moreover, since neither the spherical harmonics nor the Zernike polynomials are compact in space, evaluating a BRDF represented in terms of these functions requires computation time proportional to the total number of nonzero coefficients. Therefore, because large numbers of coefficients are necessary to avoid artifacts, evaluating BRDFs stored in terms of these basis functions is expensive.

Wavelets have been proposed as an alternative set of basis functions for BRDF representation, because they help to reduce evaluation time and storage cost. Because wavelets are localized in space, evaluating the BRDF at a particular pair of angles of incidence and reflection requires computation time proportional to the depth of the coefficient tree (i.e. logarithmic in most cases) rather than time proportional to the total number of nonzero coefficients. In addition, the spatially-localized wavelets can represent the large spikes (e.g. the specular peak) of many common BRDFs more efficiently than spherical harmonics or Zernike polynomials.

Schröder and Sweldens have proposed a basis of wavelets optimized for representing functions on a sphere [Schröder 95], although they did not actually describe an implementation of storing complete BRDFs using these spherical wavelets. Lalonde and Fournier describe a complete implementation, using wavelets on the Nusselt embedding of the hemisphere [Lalonde 97] and a tree-based encoding of the coefficients. Their experience shows that, as expected, using wavelets results in significant BRDF compression.

### 3 Change of Variables

Decomposition into basis functions is certainly a suitable technique when it is necessary to have the ability to represent arbitrary BRDFs. All of the sets of basis functions we have described, however, share the problem of requiring large numbers of coefficients to describe even moderately specular BRDFs. The main reason for this inefficiency is the fact that a BRDF parameterized in terms of angles of incidence and reflection usually does not have “localized” change. For example, a shiny surface will have a large specular peak whose position in terms of the reflectance angle  $(\theta_o, \phi_o)$  varies rapidly as the angle of incidence is changed. Although in most cases *we* know where the specular peak will lie (i.e. mostly in the direction of ideal specular reflection), the usual basis functions do not take advantage of this fact, and therefore cannot represent the peak efficiently. Similarly, the standard methods in general do not require any less storage if the BRDF is isotropic. It is possible to design a basis specialized for storing only isotropic BRDFs, but it is up to the *user* to specify that these, not the generic basis functions, should be used.

We propose an approach for making decomposition of BRDFs in terms of such functions more efficient while retaining the ability to represent (in the limit) arbitrary BRDFs. The idea was inspired by considering the possibility of using basis functions that are more “tuned” to representing common BRDFs. Instead of doing this, however, we propose the opposite approach: transforming BRDFs such that they can be represented more efficiently using the traditional basis functions.

The transformation we propose is a change of variables that causes some features of common BRDFs to lie along the new coordinate axes. Specifically, we propose parameterizing the BRDF in terms of the halfway vector (i.e. the vector halfway between the incoming and reflected rays) and a “difference” vector, which is just the incident ray in a frame of reference in which the halfway vector is at the north pole (see Figure 2). That is, instead of representing the BRDF as  $\beta = \beta(\theta_i, \phi_i, \theta_o, \phi_o)$ , we regard the BRDF as a function of the halfangle and difference angle:  $\beta = \beta(\theta_h, \phi_h, \theta_d, \phi_d)$ , where

$$\vec{n} = \text{Surface normal}, \quad (1)$$

$$\vec{t} = \text{Surface tangent (i.e. orientation of an anisotropic surface)}, \quad (2)$$

$$\vec{b} = \text{Surface binormal (i.e. } n \times t), \quad (3)$$

$$\vec{\omega}_i = \text{sph}(\theta_i, \phi_i), \text{ where } \theta_i \text{ is measured relative to } \vec{n} \text{ and } \phi_i \text{ is relative to } t, \quad (4)$$

$$\vec{\omega}_o = \text{sph}(\theta_o, \phi_o), \quad (5)$$


$$\vec{h} = \text{sph}(\theta_h, \phi_h), \quad (6)$$

$$= \frac{\vec{\omega}_i + \vec{\omega}_o}{\|\vec{\omega}_i + \vec{\omega}_o\|}, \quad (7)$$

$$\vec{d} = \text{sph}(\theta_d, \phi_d), \quad (8)$$

$$= \text{rot}_{\vec{b}, -\theta_h} \text{rot}_{\vec{n}, -\phi_h} \vec{\omega}_i. \quad (9)$$

Note that  $(\theta_h, \phi_h)$  are the spherical coordinates of the halfway vector in the  $\vec{t}-\vec{n}-\vec{b}$  frame. The two rotations in equation 9 bring the halfangle  $\vec{h}$  to the north pole, and  $(\theta_d, \phi_d)$  are the spherical coordinates of the incident ray in this transformed frame.



**Figure 2:** Proposed reparameterization of BRDFs. Instead of treating the BRDF as a function of  $(\theta_i, \phi_i)$  and  $(\theta_o, \phi_o)$ , as shown on the left, we consider it to be a function of the halfangle  $(\theta_h, \phi_h)$  and a difference angle  $(\theta_d, \phi_d)$ , as shown on the right. The vectors marked  $\vec{n}$  and  $\vec{t}$  are the surface normal and tangent, respectively.

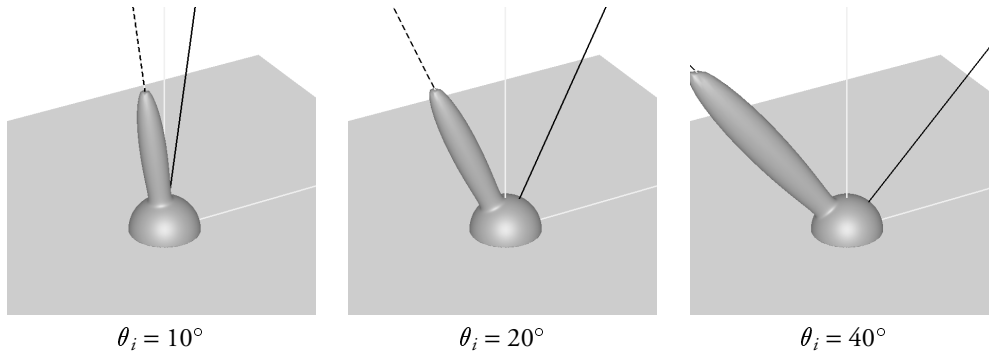
### 3.1 Properties of BRDFs in the New Coordinates

From the point of view of BRDF representation, the main effect of the proposed change of variables is that it aligns the features of common BRDFs (such as specular and retroreflective peaks) with the new coordinate axes. Thus, representing most BRDFs in terms of basis functions in these new coordinates should require a smaller number of nonzero coefficients than would be required in the untransformed coordinates for equivalent accuracy. The reason for this is that in the new coordinates the BRDFs show strong dependence on each axis individually, but show only weak dependence on combinations of axes. That is, while a BRDF might depend on  $\theta_h$  or  $\theta_d$ , most common BRDFs will not have a dependence that is some complex function of both  $\theta_h$  and  $\theta_d$ . Therefore, coefficients that correspond to terms that are high-frequency in both  $\theta_h$  and  $\theta_d$  will, for most BRDFs, be small. The only large coefficients should be the ones that correspond to variation in only one axis.

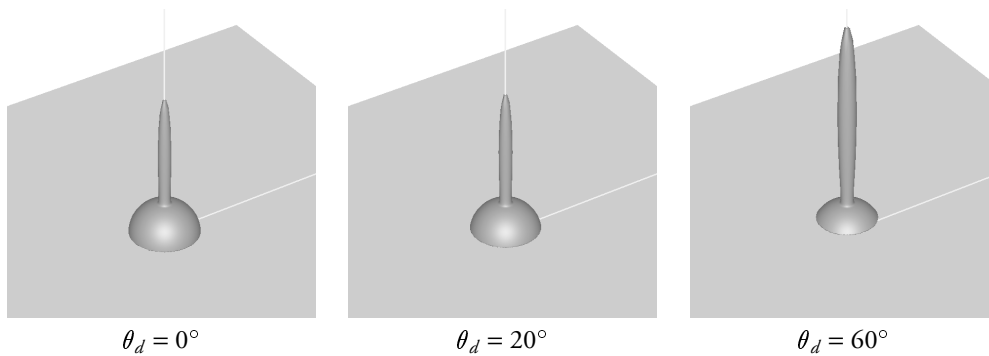
Let us now examine how certain BRDFs appear in the transformed coordinates. First, we note that isotropic BRDFs are independent of  $\phi_h$  in this coordinate system. This means that an isotropic BRDF will have basis function coefficients equal to zero for all basis functions that vary with  $\phi_h$ . Therefore, we have automatically reduced the number of nonzero coefficients to a three-dimensional subset of the four-dimensional space. This contrasts with the standard coordinates, where the entire four-dimensional space will be populated with nonzero coefficients even in the case of an isotropic BRDF.

A second property of the new coordinates is that the angles of incidence and reflection become much more symmetric. In particular, the condition of Helmholtz reciprocity becomes a simple symmetry under  $\phi_d \rightarrow \phi_d + \pi$ . It is therefore easy to enforce reciprocity in any representation based on this change of variables.

Ideal specular and near-ideal specular peaks are transformed by the change of variables to lie mostly along the  $\theta_h$  axis. An ideal specular peak is represented as a delta function of  $\theta_h$ , and is completely independent of the other three variables. Similarly, a simple BRDF such as Blinn's variation [Blinn 77] on Phong's model [Phong 75] is also a function of only  $\theta_h$ . In general, any BRDF that depends only on  $(\vec{n} \cdot \vec{h})$  is independent of three of the four variables in the transformed space. In terms of representation, this means that only a one-dimensional subset of the four-dimensional space of coefficients will be nonzero for such a BRDF.




The Cook-Torrance-Sparrow BRDF seen as a function of  $(\theta_o, \phi_o)$ , for various values of  $(\theta_i, \phi_i)$ .  
 Note that the position of the peak in space varies considerably.




The Cook-Torrance-Sparrow BRDF seen as a function of  $(\theta_h, \phi_h)$ , for various values of  $(\theta_d, \phi_d)$ .  
 Note that although the size of the peak changes (as predicted by the Fresnel term), the position and shape of the peak remain constant. The BRDF is therefore approximated very closely by a function of the form  $\beta = \beta_1(\theta_h)\beta_2(\theta_d)$ , which means that only a small number of basis function coefficients will be nonzero.

**Figure 3:** Cook-Torrance-Sparrow BRDF in standard and transformed coordinates.




The Hapke/Lommel-Seeliger BRDF seen as a function of  $(\theta_o, \phi_o)$ , for various values of  $(\theta_i, \phi_i)$ .

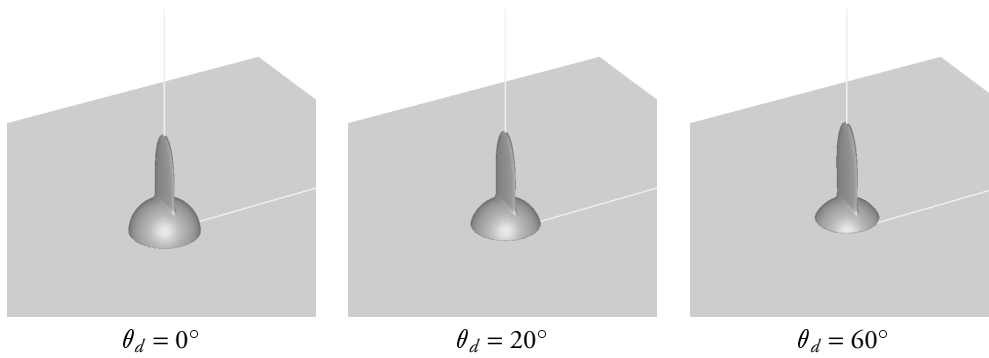


The Hapke/Lommel-Seeliger BRDF seen as a function of  $(\theta_d, \phi_d)$ , for various values of  $(\theta_b, \phi_b)$ . Compare this to Figure 3, where we plot slices of constant  $\theta_d$ . In that figure, it is the specular peak that stays relatively stationary; here it is the retroreflective peak.

**Figure 4:** Hapke/Lommel-Seeliger BRDF in standard and transformed coordinates.



Ward's elliptical Gaussian BRDF seen as a function of  $(\theta_o, \phi_o)$ , for various values of  $(\theta_i, \phi_i)$ .



Ward's elliptical Gaussian BRDF seen as a function of  $(\theta_h, \phi_h)$ , for various values of  $(\theta_d, \phi_d)$ . The BRDF is very closely approximated by a function of the form  $\beta = \beta_1(\theta_h)\beta_2(\phi_h)$ .

**Figure 5:** Elliptical Gaussian BRDF in standard and transformed coordinates.



A further generalization of the above is that any analytic model that starts with an isotropic microfacet distribution and ignores masking, shadowing, and interreflectance will be a separable function of only  $\theta_h$  and  $\theta_d$  in the transformed coordinates. To prove this, we observe that the halfangle is the direction in which a microfacet’s normal must be oriented such that light will reflect from the given direction of incidence in the given angle of reflection. Therefore, the number of microfacets pointed in the “correct” direction must be some function of  $\theta_h$ . This amount must then be corrected depending on the orientation of the incident ray relative to the microfacet (e.g. to include a Fresnel term). This term will be a function of only  $\theta_d$ , and therefore the resultant BRDF must be of the form  $\beta = \beta_1(\theta_h)\beta_2(\theta_d)$ .

More complex BRDFs will include effects such as masking, shadowing, and interreflectance, and therefore will not have completely separable representations. However, in many cases the decomposition in terms of our proposed coordinates will still be fairly simple, and may be nearly separable (see Figure 3, which shows a Torrance-Sparrow BRDF). This means that it is inexpensive to store accurate approximations even for these more realistic BRDFs.

Retroreflective peaks are transformed by our change of variables to be functions of only  $\theta_d$ : an ideal retroreflective peak would be a delta function around  $\theta_d = 0$ , and a more diffuse peak is represented by some smoother function of  $\theta_d$ . Again, in many cases more complicated retroreflective BRDFs come close to being functions of only one variable and have only a weak dependence on the other three. For example, Figure 4 shows the primarily retroreflective Hapke/Lommel-Seeliger BRDF.

Finally, let us consider one particular kind of anisotropy, namely the anisotropy associated with BRDFs that have elliptical specular peaks. These are the kinds of BRDFs predicted by, for example, Ward’s elliptical Gaussian BRDF [Ward 92], and are commonly seen on surfaces such as brushed metals. We observe that in transformed coordinates these BRDFs are again largely separable, and their main features can be represented as the product of some function of  $\theta_h$  (representing the shape of the specular peak) and some function of  $\phi_h$  (representing the anisotropic variation). Figure 5 shows Ward’s BRDF in our transformed coordinates.

### 3.2 Results

We now present some examples of how the proposed change of variables can reduce the cost required to store BRDFs. First, let us look at some data obtained from analytic models (see Figure 6). Three analytic BRDFs were randomly sampled, then least-squares fits to cubic wavelet basis functions were performed in both the standard and transformed coordinates. The fits were done to the logarithm of the BRDF rather than the BRDF itself, which had the effect of minimizing relative, rather than absolute, error. This is a more appropriate error metric for perceptual differences, and so the results should be more applicable to most applications than if the absolute error had been minimized.

Figure 6 shows the required storage for data derived from the analytic BRDFs. The first is the Torrance-Sparrow BRDF<sup>1</sup>, a standard glossy-surface model. The second is the Hapke/Lommel-Seeliger BRDF<sup>2</sup>, which describes a (mostly retroreflective) dusty surface. The third is Ward’s elliptical Gaussian model<sup>3</sup>, which describes reflection from an

---

<sup>1</sup>The parameters used were:  $m = 0.2$ ,  $\eta = 2.0$ ,  $\kappa_s = 0.5$ , and  $\kappa_d = 0.5$ .

<sup>2</sup>The parameters used were:  $g = 0.6$  and forward scattering coefficient = 0.1.

<sup>3</sup>The parameters used were:  $\alpha_x = 0.2$ ,  $\alpha_y = 0.5$ ,  $\kappa_s = 0.2$ , and  $\kappa_d = 0.5$ .

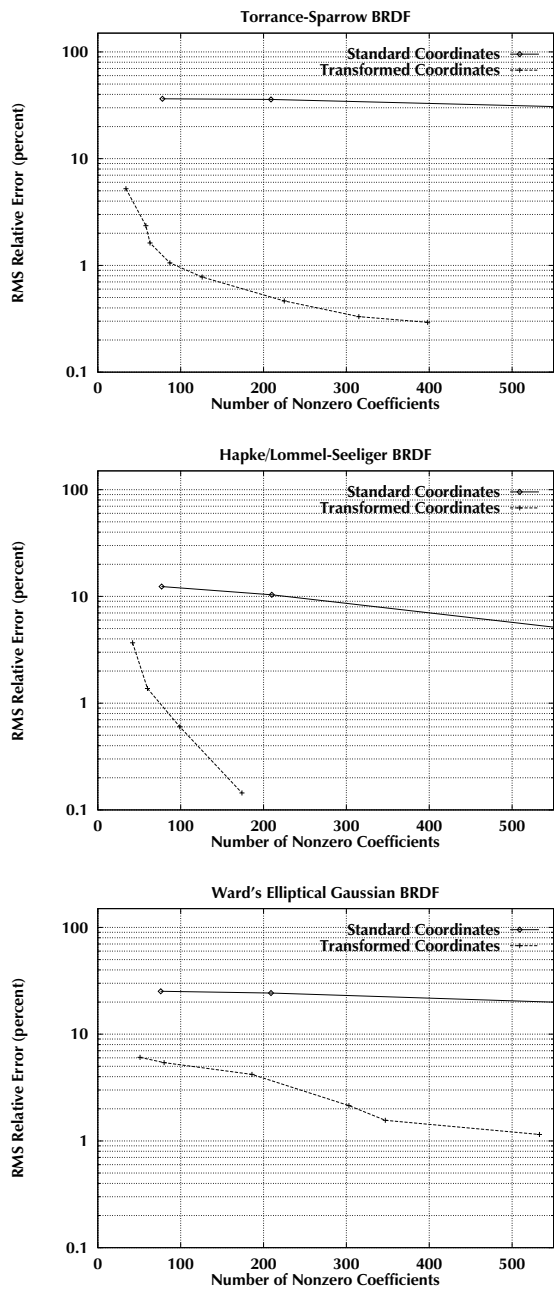
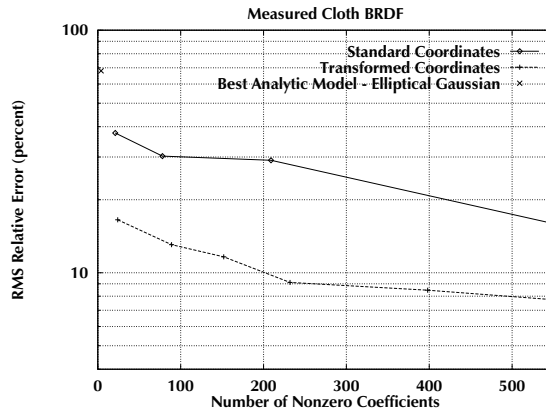


Figure 6: Representation error as a function of the number of nonnegligible wavelet coefficients in standard and transformed coordinates.



**Figure 7:** Results of fitting measured cloth BRDF to basis functions. For comparison, the best analytic result is also presented. Note that because the analytic model has a fixed number of parameters (namely, 4), it is plotted as a single point.

anisotropic rough surface. We see that in all cases performing the change of variables reduced storage requirements. Note that all of these BRDFs contained reasonably narrow (specular or retroreflective) peaks. We would expect somewhat smaller storage savings for more diffuse BRDFs.

Figure 7 shows the results obtained for the measured BRDF of a cloth sample. The cloth sample was slightly anisotropic, and exhibited markedly increased reflection towards grazing. For purposes of comparison, we also show a fit to an analytic model that best matches our data.

The cloth BRDF measurements were taken with a four-axis computer-controlled camera gantry that lets a sample be illuminated and photographed from any direction, except for an occlusion that prevents measurements within  $3^\circ$  of retroreflection. Because the data were used without any smoothing, there was about 5% variation among closely-spaced samples because of camera noise and, more importantly, surface irregularities. Therefore, we did not expect to be able to obtain errors below about 5% (using very large numbers of coefficients to try to achieve lower errors would have led to overfitting).

## 4 Conclusions and Future Work

We have presented an idea for transforming BRDFs such that they can be represented efficiently by standard classes of basis functions. The transformation:

- Reduces the number of basis coefficients necessary to represent a broad range of BRDFs.
- Allows significant savings when storing isotropic BRDFs, as compared to anisotropic ones.
- Exhibits greater savings for increasingly specular BRDFs, since it does a good job of aligning their high-frequency components (namely specular and retroreflective peaks) with the transformed coordinate axes.

Although we have concentrated on applying this transformation to fitting a BRDF to linear combinations of basis functions, it should be possible to apply the same idea to

nonlinear fits to arbitrary sets of functions. This should increase compression even further, at the cost of increased running time and possible numerical instability in the nonlinear fitting routines.

In addition, the possible space of transformations on BRDFs should be explored more widely. We have presented only one possible transformation, and surfaces displaying different kinds of phenomena, especially highly anisotropic surfaces, might benefit from different transformations. Most BRDF research in computer graphics has concentrated on diffuse and glossy surfaces, since analytic models for these were the easiest to obtain. As BRDF measurement becomes more practical and common, we plan to measure a greater variety of complicated, "exotic" BRDFs, and explore their properties, including their behavior under this change of variables.

## Acknowledgments

I would like to thank my advisor Marc Levoy for his helpful comments on the ideas presented here. I also thank the National Science Foundation for their financial support. Our BRDF-acquisition research is funded by Interval Research and Honda.

## References

- [Blinn 77] Blinn, J. "Models of Light Reflection for Computer Synthesized Pictures," *Proc. Siggraph*, 1977.
- [Cook 81] Cook, R. and Torrance, K. "A Reflectance Model for Computer Graphics," *ACM Computer Graphics*, Vol. 15, No. 4, 1981.
- [Cook 84] Cook, R. "Shade Trees," *Proc. Siggraph*, 1984.
- [Hanrahan 90] Hanrahan, P. and Lawson, J. "A Language for Shading and Lighting Calculations," *Proc. Siggraph*, 1990.
- [Hapke 63] Hapke, B. "A Theoretical Photometric Function for the Lunar Surface," *Journal of Geophysical Research*, Vol. 68, No. 15, 1963.
- [He 91] He, X., Torrance, K., Sillion, F., and Greenberg, D. "A Comprehensive Physical Model for Light Reflection," *Proc. Siggraph*, 1991.
- [Koenderink 96] Koenderink, J., van Doorn, A., and Stavridi, M. "Bidirectional Reflection Distribution Function Expressed in Terms of Surface Scattering Modes," *Proc. European Conference on Computer Vision*, 1996.
- [Lafortune 97] Lafortune, E., Foo, S. C., Torrance, K., and Greenberg, D. "Non-Linear Approximation of Reflectance Functions," *Proc. Siggraph*, 1997.
- [Lalonde 97] Lalonde, P. and Fournier, A. "Filtered Local Shading in the Wavelet Domain," *Proc. Eurographics Workshop on Rendering*, 1997.
- [Minnaert 41] Minnaert, M. "The Reciprocity Principle in Lunar Photometry," *Astrophysical Journal*, Vol. 93, 1941.
- [Oren 94] Oren, M. and Nayar, S. "Generalization of Lambert's Reflectance Model," *Proc. Siggraph*, 1994.
- [Phong 75] Phong, B. T. "Illumination for Computer-Generated Pictures," *Communications of the ACM*, Vol. 18, No. 8, June 1975.
- [Schröder 95] Schröder, P. and Sweldens, W. "Spherical Wavelets: Efficiently Representing Functions on the Sphere," *Proc. Siggraph*, 1995.
- [Ward 92] Ward, G. "Measuring and Modeling Anisotropic Reflection," *Proc. Siggraph*, 1992.
- [Westin 92] Westin, S., Arvo, J., and Torrance, K. "Predicting Reflectance Functions from Complex Surfaces", *Proc. Siggraph*, 1992.